intel

How To Select A Microcomputer Bus Architecture

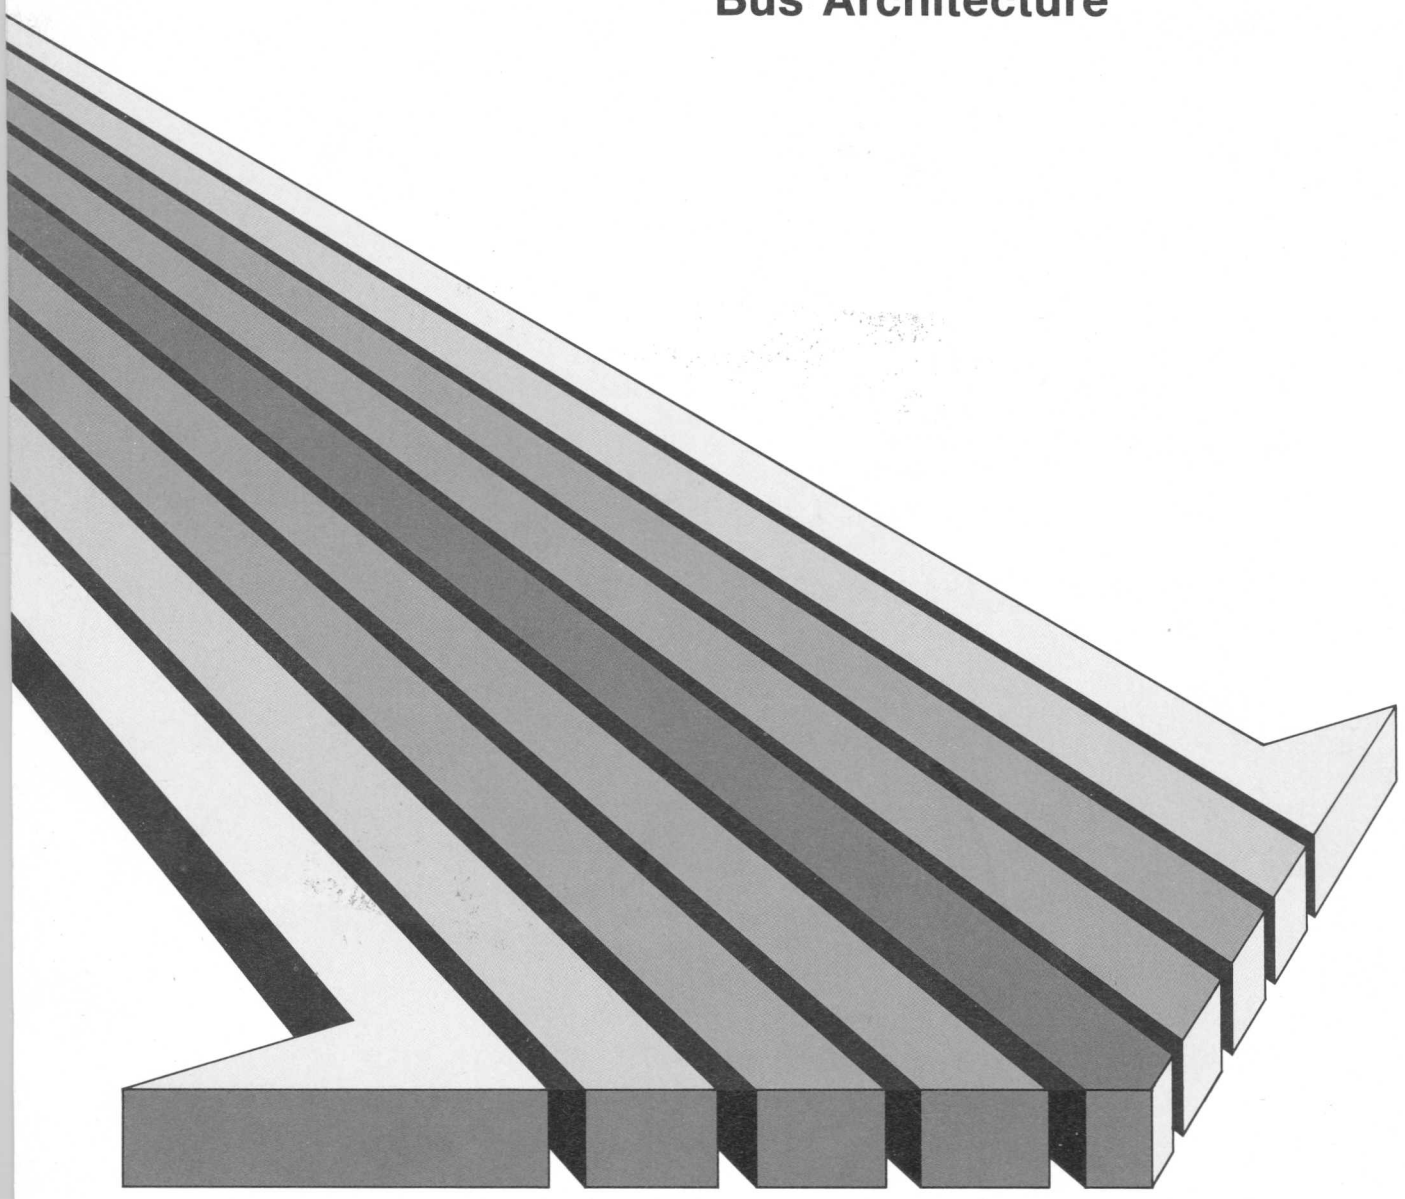

Table of Contents

230727-001

HERE'S TO THE 200 WHO BOARDED OUR BUS

Multibus-product vendors are forming trade group

The association aims to hasten standardization and preserve Multibus's leadership position

The most popular data pathway for use in computers is Multibus. The worldwide market for Multibus products, such as single-board microcomputers, memory boards and other devices that hook onto the data path, today totals about $150 million. According to Frederick J. Mazanec, a Multibus consultant and president of Ironoak Co. in La Jolla, Calif., that market will climb to more than $350 million by 1985. He estimates that including software firms, there are now 170 vendors of Multibus-compatible products worldwide.

Multibus, designed by Intel Corp. in 1976, completely dominates the 8-bit market. Its influence in the 16-bit field — in which it competes head-on with Digital Equipment Corp.'s Unibus — is "growing," Mazanec asserts.

However, it faces potential strong competition from VME, a bus structure backed by Mostek Corp., Motorola Corp., the Signetics Corp. division of North American Philips Corp., and Thomson-CSF. VME sponsors say that bus, which is popular in Europe and gaining acceptance in the United States, also will support 16- and 32-bit hardware in single- and multiple-processor systems.

To maintain Multibus's competitive edge, a Multibus Manufacturers Group (MMG) is in the works. According to Nicholas Robertshaw, engineering manager for the Plessey Microsystems division of Plessey Inc. and chairman of the MMG steering committee, the group's constitution probably will be formalized this month at Electro in New York City or next month at the National Computer Conference in Anaheim, Calif. Robertshaw says the group will help "unify the market image for the Multibus."

Not surprisingly, Intel is assisting Robertshaw in the formation of the unit because it will give the Santa Clara, Calif., company another opportunity to promote its products, which include a line of 16-bit single-board Microbus computers introduced in January.

While new products are always eye-openers, potential MMG members probably are more interested in Intel's long-term Multibus plans, especially Multibus II. However, though Intel has acknowledged publicly that it is working on Multibus II, details of the new bus structure are yet to be released. □

Multibus dominates the worldwide market for bus products

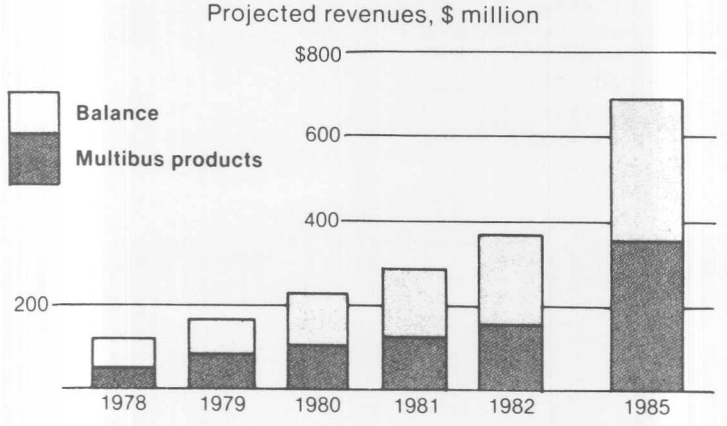

Source: Ironoak Co.

Intel Creates Multibus Group

HILLSBORO, ORE. -- Intel Corp. recently formed the Multibus Manufacturers Group, independent of the company, to help promote the product's use by other manufacturers.

Named chairman of the group's steering committee was president Fred Mazanec of La Jolla, Calif.-based Ironoak Co. The steering committee is comprised of representatives of 20 companies currently marketing Multibus-based products.

Intel estimated the market for single-board computers last year to be about $375 million in sales, approximately one half of which were Multibus compatible.

By 1985, the value of revenues is expected to grow to almost $800 million, according to Intel, with Multibus-compatible products expected to maintain approximately 50% of the market share. Currently, about 150 firms supply more than 1000 different Multibus-based products.

The group plans to do joint promotion of the Multibus in the hope that the product's market share will grow, according Mr. Mazanec.

"The reason for the group is the manufacturers hope to promote the Multibus for the good of the members," he explained. "That way, everyone's piece of the pie will grow."

Other members of the steering committee include: Micro Industries, Advanced Micro Devices, Plessey Microsystems, Meticomp, Central Data Corp., NEC Information Systems Inc., Vertex Corp., Omnibyte Corp., Interphase Corp., Microbar Systems Inc., Systech Corp. and Electronic Solutions.

Another 50 or 60 Multibus-product makers are interested in joining the group, Mr. Mazanec added.

MACHINE DESIGN

Multibus manufacturing group formed

A Multibus manufacturers group has been formed to provide oversight and guidance to Multibus-related issues. The Multibus is a widely-used format standard for board-level microcomputer products. Frederick Mazanec, president of Ironoak Co. of La Jolla, Calif., has been named chairman of the new Multibus group steering committee. The steering committee consists of representatives of 20 companies currently marketing Multibus-based products.

Formation of the manufacturers group reflects the rapid growth of Multibus-compatible products. In 1982, for example, these products accounted for approximately one-half of the estimated $375 million sales of single board computers. Currently, approximately 150 firms supply over 1,000 Multibus based products. In late 1982, the Multibus configuration was adopted as an industry standard by the Institute of Electrical Engineers (IEEE).

DATAMATION

Now that Intel's Multibus has been adopted by the IEEE as an industry standard microcomputer bus, it seems appropriate that the Hillsboro, Ore., vendor has established a Multibus Manufacturers Group to provide oversight and guidance on related industry issues. Over 150 firms supply Multibus products.

DESIGNER'S GUIDE SERIES:

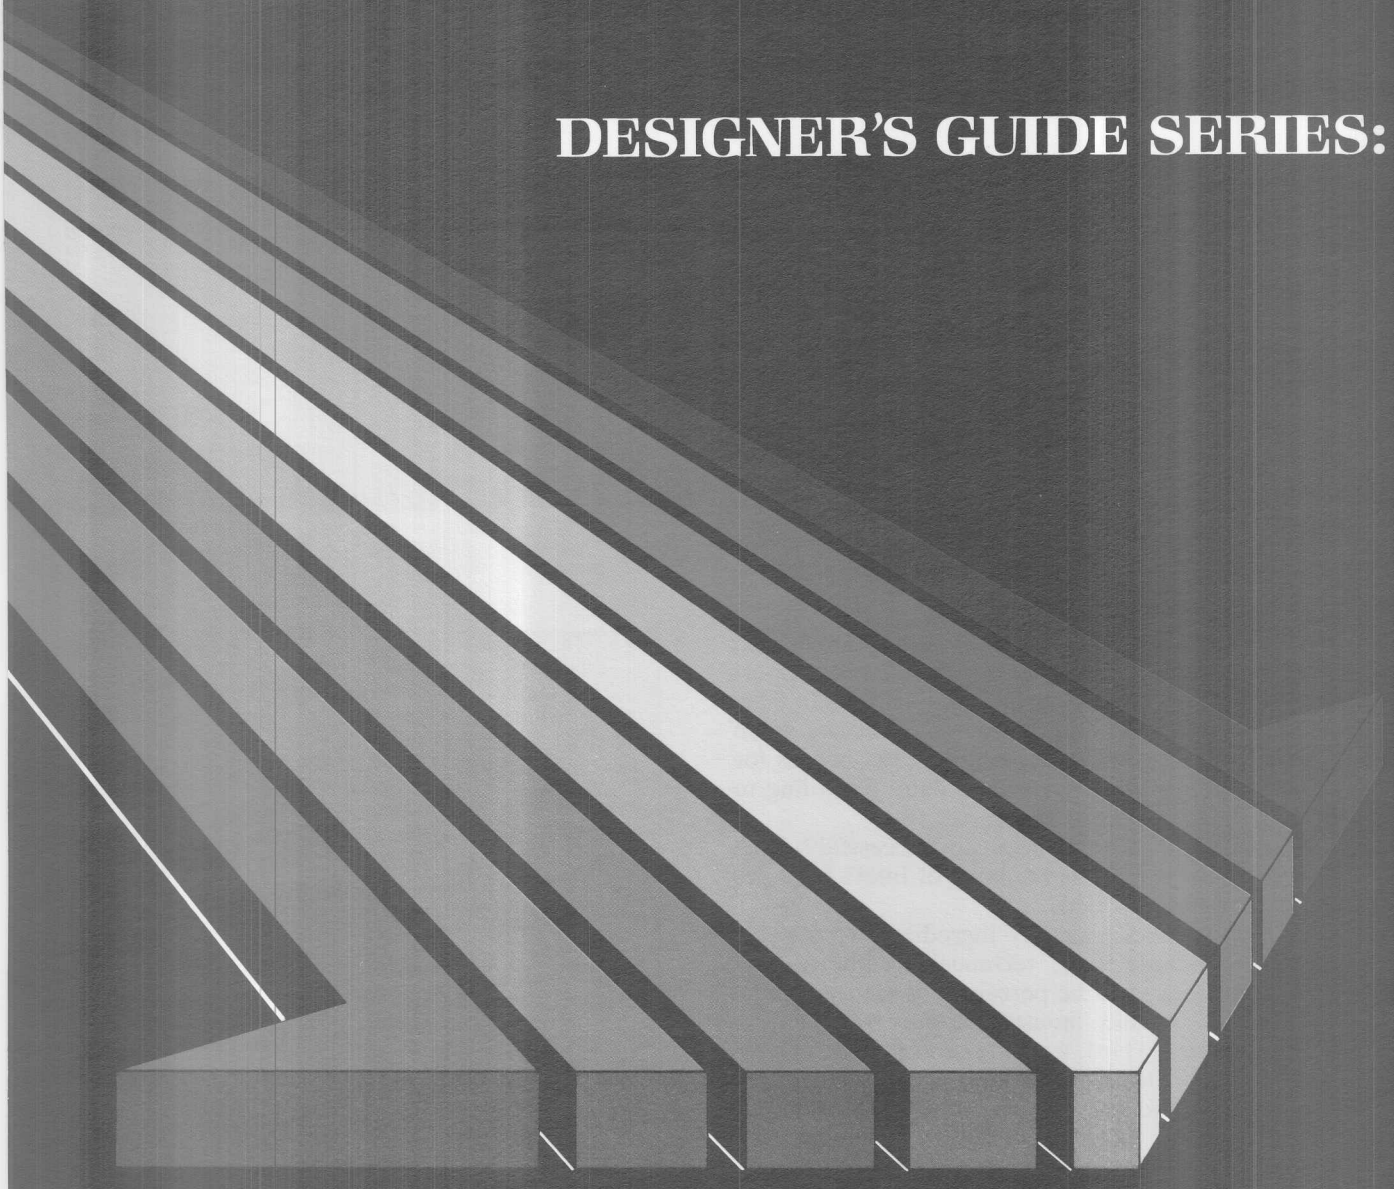

THE MULTIBUS

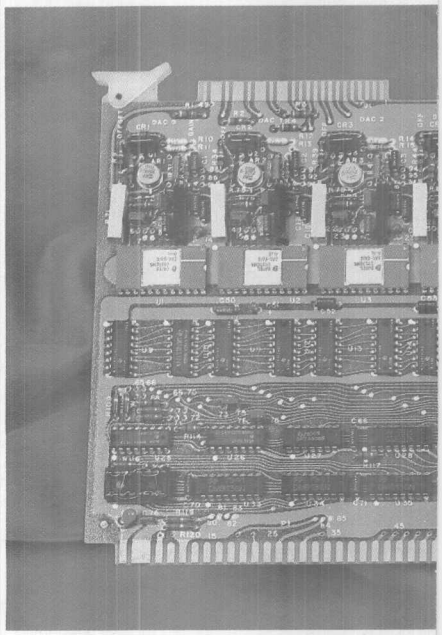

Multibus: Evolving To Meet New System Demands

by Dave Wilson, Senior Technical Editor

With over 800 board level products available from over 150 vendors, the Multibus has evolved to support three generations of VLSI technology—the 8080, 8086 and now the 286. As a testimony to its success as an industry standard, over half a million CPU boards for the Multibus have been shipped to date, according to Intel Corp.

Why has the Multibus been so successful? Rich Bader, Strategic Marketing Manager of Intel's Modules Operations, has some ideas.

"We think it takes four key ingredients to make a successful Standard. First, technical credibility—the proposed standard must be perceived as viably filling a technical need. Second, industry support by a number of market leaders. Third, a strong central point of control over the specification is needed to ensure compatibility with the Standard. With the approval of Multibus spec by the IEEE, they will now play that role. And fourth, the demonstration of the ability to evolve the standard to satisfy new requirements and do so in a compatible fashion."

Multibus History

The Multibus architecture has a long history of evolution. Previously, Intel incorporated the iSBX bus for on-board I/O expansion, dual port memory, established enhanced multiprocessing support with the intelligent slave interface, and a multiprocessing soft-

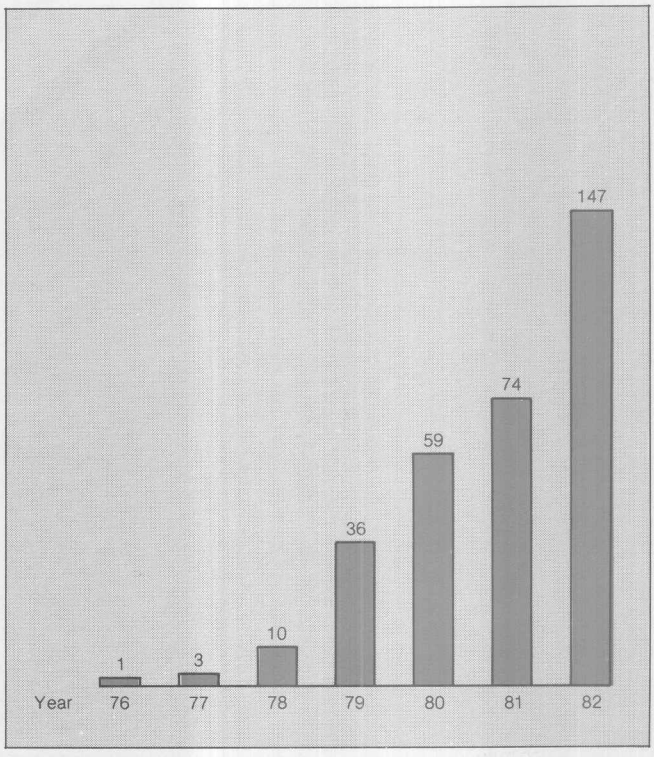

Figure 1: Multibus vendor growth.

Cover design courtesy Intel Corp.

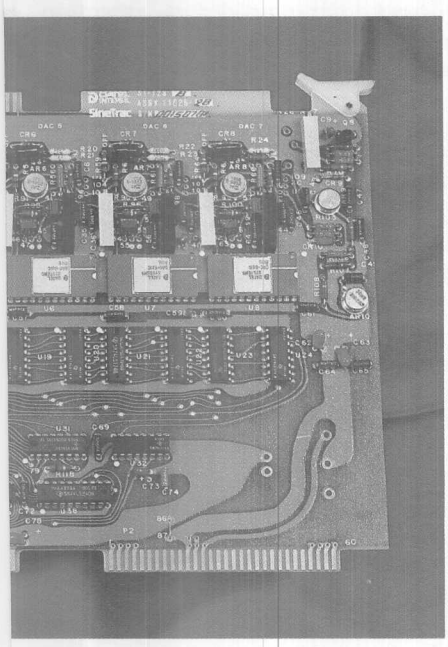

Photo courtesy Datel-Intersil

The Multibus supports three generations of VLSI technology—the 8080, 8086, and now the 286. The role of the Multibus architecture will evolve to meet the demand for 32- bit systems.

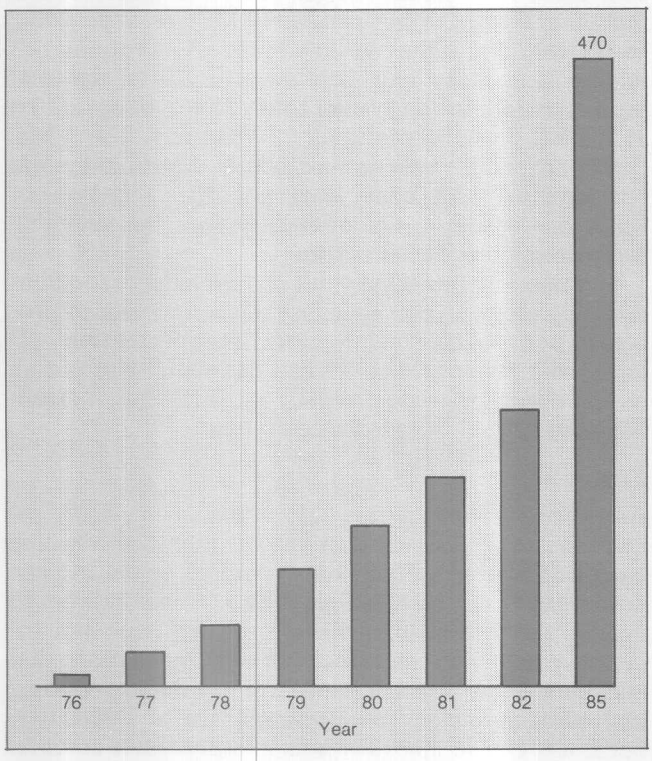

Figure 2: Multibus product revenue ($M).

ware package called MMX 800 (Multibus Message Exchange). All are part of what Intel claims is a consistent architecture for high performance, cost-effective microcomputer systems. Intel has been busy over the last year continuing this evolution, adding two additional bus structures: the iLBX Bus, a high speed execution bus; and the Multichannel I/O Bus, for very high-speed I/O. According to Rich Bader, these structures support a consistent architecture and are something unique to the Multibus system.

"Just a bus structure is no longer sufficient to take maximum advantage of the new VLSI technology. The Multibus represents an architectural concept with all of the features needed to implement it. Our customers have told us the architectural model is what they needed most, and we agree it's the key to long-term success." Multibus vendor growth is indicated in **Figure 1**; Multibus product revenue is illustrated in **Figure 2**.

RAM Requirements

As the designer builds high performance systems that include multiple SBCs, he finds that it is impossible to put all of his memory requirements onto the SBC; there is physically not enough space. RAM requirements are increasing at a rate faster than the RAM density of VLSI, due to processors such as the 286 and the 68000 that can address large memory spaces.

Intel has recognized that the architecture of a sys-

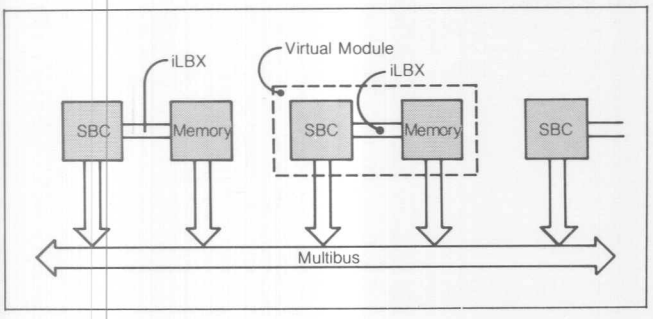

Figure 3: The iLBX bus.

tem must be able to support this and has introduced the iLBX, an add-on bus that allows the designer the ability to communicate between an SBC and memory card without tying up the Multibus. Typically, if the designer has been unable to put all the memory he requires on the SBC board, he has put the memory on the systems bus. This, however, causes the system to be used for execution. With the kind of processor performance that is available today, system buses cannot handle execution, DMA and interprocessor communication, and have sufficient bandwidth. For example, the Multibus bandwidth is not large enough to efficiently support multiple iAPX 286 or 68000 SBCs and memory boards. Bader claims, "VME bus users will find out very quickly they have the same problem." As a consequence, if the user intends to build big memory intensive systems, he needs a bus that can move the execution off the system bus and allow the system bus to focus on other things, such as DMA and interprocessor communication and have a separate high bandwidth bus to support the memory.

Enter The iLBX

With the iLBX bus the user can have as many virtual modules, consisting of SBC and memory as he requires **(Figure 3)**. From an architectural standpoint, having the SBC and the memory board all talking to the Multibus, the collection of boards looks like a collection of SBCs with all the memory physically on the card. A number of other companies have already introduced similar bus structures like this, e.g. Microbar. (Microbar call their aproach a dual bus.) Microbar will continue to support the dual bus, but they have announced their intention to support the iLBX bus in the future, as well. Microbar may perceive that with Intel backing the iLBX will become the next element in the industry standard architecture. Microbar's Dual Bus computer family features both the 8086 and 68000 μPs and utilizes a high speed dual bus architecture very similar to Intel's recently announced bus extention. "Intel's iLBX validates our dual bus architecture" says Rich Boberg, President of Microbar systems, "and because of the design similarities, compatibility with the iLBX will be quick and easy." Microbar will offer the iLBX for both the 8086 and 68000 based CPU boards, as well as dual-ported memory for demanding high-performance applications. Interphase, Zendex, Central Data and Plessey have also announced their intent to support the iLBX bus.

Spec It

The 796 specification covers the P1 connector on the Multibus and includes 4 address lines on the P2 to allow for 20-bits of address. The remaining lines on the P2 connector were previously declared in the spec to be reserved. The iLBX will be physically on the P2 connector giving a standard definition of a function on that connector. The iLBX will be typically used for processors talking to memory.

However, if the designer has an application where he wants to use the P2 connector for other functions, there is no reason why this cannot be done. For example, if he has an A/D converter board he wishes to interface to the Multibus and does not need the iLBX bus, there is no reason to dedicate the pins for that function. The iLBX bus runs at 19 MB/sec when doing 16-bit transfers. If the user has 3 processors running, all talking to memory over their own iLBX bus, the total bus bandwidth in the system totals to 67MB/sec. (19 x 3) iLBXS plus the Multibus (10MB/sec.). The VME Bus is not significantly faster than the Multibus system bus structure and VME does not offer this type of iLBX bus—they do not offer a systems bus architecture, simply one bus structure.

In summary, bus bandwidth requirements are about doubling every 2 years, because of the increase in processor performance. With more multiprocessing systems, with faster processors and the inability to put all required memory onto an SBC, the iLBX bus is essential to build high performance systems.

Faster I/O

As processors are getting faster and faster, users require faster I/O as well. Just as execution can consume bandwidth on the systems bus, very high speed I/O can do the same. As the Multibus architecture continues to evolve, the system bus will be used for what can generically be called interprocessor communication; that may be a smart disk talking to a smart data processing unit, so it may have some resemblance of DMA I/O, but it will evolve more towards interprocessor communication as all of the boards become smart and they preprocess data that is being passed within the system.

With the architectural desire to devote the Multibus system bus to interprocessor communication, and in order to satisfy the need for higher performance I/O, Intel needed a very high speed bus for I/O transfers. Their customer base stated requirements for a fast bus that would allow them to run the bus outside the card cage (all system buses are card cage restricted), from box to box to a data acquisition system, for example. The customers also wanted a solution to the architectural problem. The Intel solution (**Figure 4**) was to create yet another bus structure called the Multichannel bus. In order to provide more buffer memory than implemented on the controller, the bus can also talk directly to the backside of additional memory boards. This allows the designer to move data into memory rapidly and move the data around at his leisure on the Multibus.

The DMA

The DMA controller, (**Figure 4**), dubbed the iSBC 589, can have control over the access of data, and according to the Multichannel specification, the user may also build on smart I/O devices to be able to do the transfer directly. **Figure 5** shows a system configuration using a Multichannel and iLBX interface for additional buffer memory. The SBC can access the system memory over the iLBX, and the Multichannel device can transfer information thru the 580 into the memory. It is left up to the designer to decide if he requires the 580 to transfer data over an iLBX bus that is common or separate from the SBC (see dotted lines), to establish the proper bus bandwidth balance. The 580 converts the protocol of the Multichannel bus to the protocol of the iLBX bus.

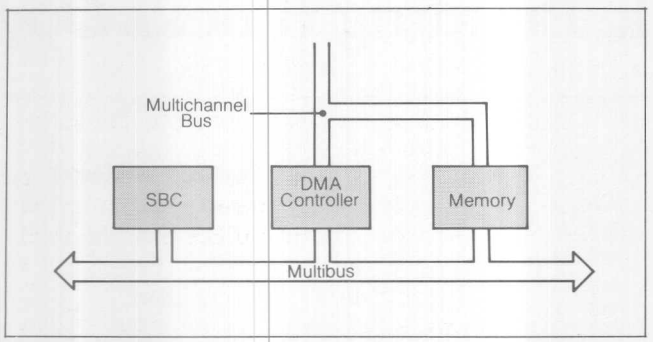

Figure 4: The Multichannel bus.

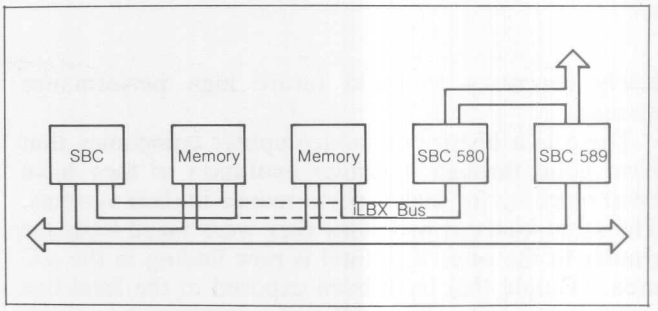

Figure 5: Multichannel iLBX interface.

The multichannel bus is capable of tranferring data at 8 MB/sec. Using a block transfer protocol, it can transfer up to 15 meters (50 ft.) supporting up to 16 devices. The Multichannel protocol is similar to the IEEE 488 bus, in that it has talker/listener and controller functions—the difference is that it is significantly faster.

What About VME?

Many rumblings were heard at Comdex Fall '82 about the VME bus and its support of 32-bit machines. One manufacturer even said he wished Intel would get a move on and support it themselves. Intel's Rich Bader stated:

"Motorola has been trying to position the VME bus as a high performance bus. But if you look at the spec, it is only marginally faster than Multibus—we looked at system bus throughput—it's about 20% faster. However, because of the kind of technological change we're seeing, it's clearly insufficient. Major enhancements, like the execution bus are needed to keep up with the technology. My contention is that our systems bus architecture is a clear head and shoulders performance winner over the VME bus. Our architecture is abso-

lutely necessary to build future high performance systems."

There is a history of minicomputer companies that have gone through a similar evolution as they have been reaching for higher performance in their systems. The same kinds of problems they were faced with are similar to the ones that Intel is now finding in the μC area. "People that have been exposed to the Intel bus architecture and compare it with what they have with VME say the choice is clear," continues Bader. "If they want to build those high performance systems they're going to need this architecture and the VME bus appears to be just a bus structure without an architecture. With the proven track record and products available, the Multibus is the obvious business choice as well."

There is no execution bus within the VME spec or any of the products that exist on the VME bus, and there is no high performance I/O bus that is being offered. Nor does there seem to be room to add them on the connectors. However, an I/O channel does exist, a 2MHz, 8-bit only channel that does not support DMA. Bader feels that it is analogous to the iSBX multimodules that Intel put on their SBCs to enhance the features. "It's something we did 3 years ago," he added.

What about 32-bits? "Of course Intel is developing a 32-bit evolution to the Multibus architecture, called Multibus II. Now, with our architecture more fully established, that's the next step. It will offer significantly higher performance than existing structures, and incor-

Imaging Boards – Flexible Building Blocks Ope

Imaging boards such as those manufactured by Datacube and Imaging Technology may be viewed simply as data acquisition and display systems. An analog signal from a video camera or other video source, such as a VTR, is converted to digital value by an A/D converter. The information is then stored in memory where it is available to the Multibus. Information from the memory also goes back to a D/A, where the digital data is converted to analog form for display on a video monitor.

"What makes us different is that we have intelligence on our board (an ALU) that can crunch pixels at 10^6/sec," says Bob Wang at Imaging Technology, "We're looking at the basic signal processing market with ALU capability—the fact is you can now take a camera input, store that, take another input, add it to what you've just started with, divide it by 2, and store it again."

This process of averaging data, called a running average, can reduce noise levels tremendously; a useful feature in low-light environments that use infra-red cameras.

"There are some things that ALUs are nice for," adds Dave Erikson at Datacube, "Such as subtraction of images in real time." (See *Digital Design*, October 1982). "But everything an ALU can do, a CPU/board can too—but not in real time. A complete screen may take several seconds as opposed to parts of a second. We've found that 99% of applications can be met by the CPU and they don't have to be done in real time."

"In some applications it makes no difference whether you acquire and process a frame in 16ms, if it has to be chewed up for 5s by the CPU," states Stan Karandanis, President of

Product	Resolution	Date Available
VG120	320 × 256 × 6 Bits	Now
VG121	320 × 242 × 6 Bits 320 × 484 × 6 Bits	Now
VG221	640 × 242 × 3 Bits 640 × 484 × 3 Bits	Now
VG123 2 Board Set	768 × 242 × 8 Bits 768 × 484 × 8 Bits 384 × 242 × 8 Bits 384 × 242 × 8 Bits	Now
VG124	640 × 242 × 6 Bits 640 × 484 × 6 Bits	2 months
VG131	$\binom{768}{384} \times \binom{242}{484} \times$ 8 Bits With ALU	2–4 months
VG140	512 × 512 × 4 (8) Bits 3 Port memory with very high speed vector input port	2–4 months

Table 1: Datacube's product line: Past & Future

porates support for other requirements needed in the 1980's and 1990's; like enhanced multiprocessing support, reliability and serviceability, and so on. And perhaps most importantly, you will be able to utilize existing Multibus cards in a Multibus II environment. Preserving everyone's investment in existing hardware will give Multibus II an enormous headstart on industry-wide acceptance," Bader concluded.

Multibus Structural Features

The Multibus interface is an asynchronous, multiprocessing system bus designed to perform 8-bit and 16-bit transfers between single board computers, memory and I/O expansion boards. Its interface structure consists of 24 address lines, 16 data lines, 12 control lines, 9 interrupt lines and 6 bus exchange lines. These signal lines are implemented on SBCs and a mating backplane in the form of two edge connectors resident on 6.75″ × 12″ form factor PC boards. One of the big complaints about the Multibus is that the form factor is too big. The iLBX concept gets around that by defining a set of modular I/O cards, so that the user may tailor his system to his own requirements inexpensively. The primary 86-pin P1 connector contains all Multibus signal lines except the four address extension lines. The auxiliary 60-pin P2 connector contains the four Multibus address extension lines, and reserves the remaining 56 pins for implementing the iLBX Execution Bus into Multibus system architecture.

New Markets

Datacube, "If functions such as linear to polar co-ordinate conversions, matrix shift rotate/displace and non-integer zoom could be done by the ALU, you've got yourself some features. And it can be done! It's being done by Thomson CSF but it costs $50,000. It's possible to zoom in on a portion of an image—but that takes more than just a barrel shifter and an ALU to do."

"This month we will have a VG131 (**Table 1**) at least defined," he adds. "It will be the first of our products with an ALU for hardware image integration. We're going to be talking about a single Multibus board as well, not several."

End Users Want More

"The problem with an end user is that he always wants more," continues Bob Wang. "He sees the capability of image subtraction and now he wants to measure the size of the subtracted image. That's a completely different class of problem." If, for example, a user wants to take 2 images, and subtract them, he can use one of IT's driver packages, supplied on a floppy disk and written in μP code. "While most people are happy with that right now, we will definitely have to head towards a much better support as far as software is concerned, in order to make the product more effective," says Wang.

Applications

Although "These products aren't going into the world's most sophisticated image processing systems," according to Stan Karandanis, application areas for imaging boards are broad and diverse, and they range from teleconferencing to radar and supermarket check-out counters.

Today, Datacube is currently interfacing with a customer who wants to dispense with bar codes on products and the necessity for a laser scanner at check-out counters. "We can go back to just a label that the customer can read—it will be waved over a solid-state sensor that inputs the item and price into the cash registers regardless of rotation. Machine and human can have a common language and that's important," says Karandanis. "$1,000 at every station isn't unreasonable at all, and that's existing technology," adds Erikson.

Another customer is looking at archiving X-rays, on videotape. Instead of having 54,000 frames of continuous movement, the tape holds 54,000 individual X-rays. "We can record an ID code recorded on the audio track (or control track) depending on the tape recorder, identify a unique frame, grab that off the tape and freeze it for a doctor to examine. Conventional media like Winchesters have a data rate integrity of one part in ten, but the fact of the matter is, when you're dealing with the human eye, you can tolerate more data error," says Karandanis, "And when you think of the data available to you on a one hour videocassette, you have 2^{10} Bytes of data for under $1,000. And the tape costs ten bucks!"

"Right now you can hook a camera to some of our boards, but the camera costs as much as the boards—and the system requires three boards." Datacube's new VG123 Board is a 2 board solution that can hook up to low-cost cameras and be compatible with low cost monitors. "The 123 will have the ability to lock into the NTSC color subcarrier so you can digitize full color signals and freeze them as frames. We're going to attack the teleconferencing market with that product," states Erikson.

	(COMPONENT SIDE)			(CIRCUIT SIDE)		
	Pin	Mnemonic	Description	Pin	Mnemonic	Description
Power Supplies	1	GND	Signal GND	2	GND	Signal GND
	3	+5	+5 VDC	4	+5	+5 VDC
	5	+5	+5 VDC	6	+5	+5 VDC
	7	+12	+12 VDC	8	+12	+12 VDC
	9	−5	−5 VDC	10	−5	−5 VDC
	11	GND	Signal GND	12	GND	Signal GND
Bus Controls	13	BCLK/	Bus Clock	14	INIT	Initialize
	15	BPRN/	Bus Priority In	16	BPRO/	Bus Priority Out
	17	BUSY/	Bus Busy	18	BREQ/	Bus Request
	19	MRDC/	Memory Read Command	20	MWTC/	Memory Write Command
	21	IORC/	I/O Read Command	22	IOWC/	IO White Command
	23	XACK/	XFER Acknowledge	24	INH1/	Inhibit 1 Disable RAM
	25	AACK/	Advance Acknowledge	26	INH2/	Inhibit 2 Disable ROM
	27		Reserved	28		Reserved
	29		Reserved	30		Reserved
	31	CCLK/	Constant Clk	32		Reserved
	33		Reserved	34		Reserved
Interrupts	35	INT6/		36	INT7/	
	37	INT4/		38	INT5/	
	39	INT2/	Parallel Interrupt Requests	40	INT3/	Parallel Interrupt Requests
	41	INT0/		42	INT1/	
Address	43	ADRE/		44	ADRF/	
	45	ADRC/		46	ADRD/	
	47	ADRA/		48	ADRB/	
	49	ADR8/	Address Bus	50	ADR9/	Address Bus
	51	ADR6/		52	ADR7/	
	53	ADR4/		54	ADR5/	
	55	ADR2/		56	ADR3/	
	57	ADR0/		58	ADR1/	
Data	59	DATE/		60	DATF/	
	61	DATC/		62	DATD/	
	63	DATA/		64	DATB/	
	65	DAT8/		66	DAT9/	
	67	DAT6/	Data Bus	68	DAT7/	Data Bus
	69	DAT4/		70	DAT5/	
	71	DAT2/		72	DAT3/	
	73	DAT0/		74	DAT1/	
Power Supplies	75	GND	Signal GND	76	GND	Signal GND
	77	−10	−10 VDC	78	−10	−10 VDC
	79	−12	−12 VDC	80	−12	−12 VDC
	81	+5	+5 VDC	82	+5	+5 VDC
	83	+5	+5 VDC	84	+5	+5 VDC
	85	GND	Signal GND	86	GND	Signal GND

Table 1: Multibus groupings and pin assignments.

Multibus Elements

The Multibus system bus supports three device categories: Master, Slave and Intelligent Slave. A bus master is any module that has the ability to control the bus. This ability is not limited to only one master device. The Multibus interface is capable of supporting multiple masters on the same system through bus exchange logic. Once access has been acquired by a master device, it has a period of exclusive control to affect data transfers through a generation of command signals and memory or I/O addresses.

A bus slave device is a module that decodes the address lines on the Multibus and acts upon the command signals from the bus masters. Slave devices are not capable of controlling the Multibus interface. The intelligent slave has the same bus interface attributes as the slave but also incorporates an on-board μP, memory and I/O allow the intelligent slave to complete onboard operations without access.

Interface/Signal Line Descriptions

The Multibus systems bus signal lines are grouped into five classes based on the functions they perform: 1) control lines 2) address and inhibit lines, 3) data lines, 4) interrupt lines, 5) bus exchange lines. **Figure 6** illustrates the implementation of these signal lines. The Multibus control lines are broken down into five subgroups: clock signals (2), commands (4), acknowledge (1), initialize (1), and lock (1). The two clock signals provide for the generation of a master clock for the system and the synchronization of bus arbitration logic. The four command lines are the communications links between the bus masters and bus slaves, specifying types of operations to be performed such as reads or writes from memory or I/O. The transfer acknowledge line is the slave's acknowledgement that a requested action of the master is complete. The initialize signal is generated to reset the entire system to a known state. The lock signal is used by an active bus

COMPONENT SIDE				SOLDER SIDE		
16-Bit Pin	Mnemonic	Signal Name		16-Bit Pin	Mnemonic	Signal Name
1	DB0	DATA LINE 0		2	DB1	DATA LINE 1
3	DB2	DATA LINE 2		4	DB3	DATA LINE 3
5	DB4	DATA LINE 4		6	DB5	DATA LINE 5
7	DB6	DATA LINE 6		8	DB7	DATA LINE 7
9	GND	GROUND		10	DB8	DATA LINE 8
11	DB9	DATA LINE 9		12	DB10	DATA LINE 10
13	DB11	DATA LINE 11		14	DB12	DATA LINE 12
15	DB13	DATA LINE 13		16	DB14	DATA LINE 14
17	DB15	DATA LINE 15		18	GND	GROUND
19	AB0	ADDRESS LINE 0		20	AB1	ADDRESS LINE 1
21	AB2	ADDRESS LINE 2		22	AB3	ADDRESS LINE 3
23	AB4	ADDRESS LINE 4		24	AB5	ADDRESS LINE 5
25	AB6	ADDRESS LINE 6		26	AB7	ADDRESS LINE 7
27	GND	GROUND		28	AB8	ADDRESS LINE 8
29	AB9	ADDRESS LINE 9		30	AB10	ADDRESS LINE 10
31	AB11	ADDRESS LINE 11		32	AB12	ADDRESS LINE 12
33	AB13	ADDRESS LINE 13		34	AB14	ADDRESS LINE 14
35	AB15	ADDRESS LINE 15		36	GND	GROUND
37	AB16	ADDRESS LINE 16		38	AB17	ADDRESS LINE 17
39	AB18	ADDRESS LINE 18		40	AB19	ADDRESS LINE 19
41	AB20	ADDRESS LINE 20		42	AB21	ADDRESS LINE 21
43	AB22	ADDRESS LINE 22		44	AB23	ADDRESS LINE 23
45	GND	GROUND		46	ACK*	SLAVE ACKNOWLEDGE
47	BHEN	BYTE HIGH ENABLE		48	R/W	READ NOT WRITE
49	ASTB*	ADDRESS STROBE		50	DSTB*	DATA STROBE
51	SMRQ*	SECONDARY MASTER REQUEST		52	SMACK*	SECONDARY MASTER ACKNOWLEDGE
53	LOCK*	ACCESS LOCK		54	GND	GROUND
55	ADR22*	MULTIBUS® ADDRESS EXTENSION LINE 22		56	ADR23*	MULTIBUS ADDRESS EXTENSION LINE 23
57	ADR20*	MULTIBUS® ADDRESS EXTENSION LINE 20		58	ADR21*	MULTIBUS ADDRESS EXTENSION LINE 21
59	RES	RESERVED		60	TPAR*	TRANSFER PARITY

Table 2: Pinouts for the iLBX bus.

master to lock dual ported for mutual exclusion. The address and inhibit lines are made up of 24 address lines, two inhibit lines and one byte control line. The 24 address lines are signal to carry the address of the memory location or the I/O that is being referenced. These 24 lines allow a maximum of 16 Mbytes to be accessed.

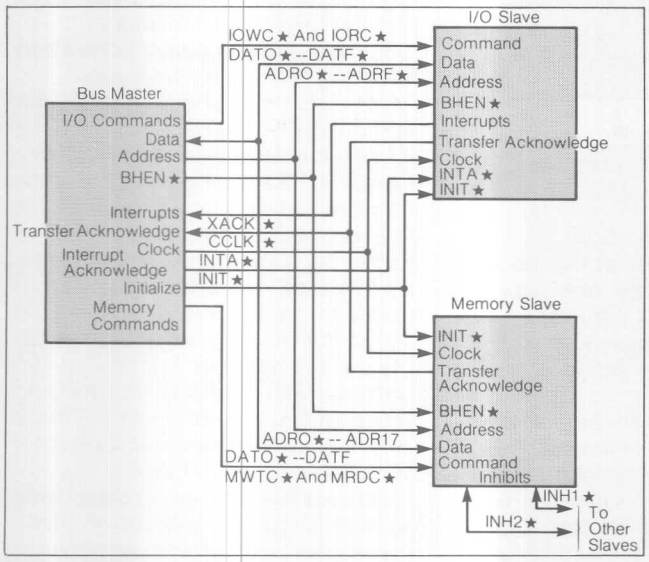

Figure 6: 796 Bus Interface Lines.

When addressing an I/O device, sixteen address lines are used to address a maximum of 64K devices. The two inhibit lines are used to allow different types of memory (RAM, ROM, etc.) having the same memory address to be accessed in a preferred priority arrangement. The byte control is used to select the upper byte of a 16-bit word in systems incorporating 16-bit memory and I/O modules. The Multibus interface supports sixteen bi-directional data lines to transmit information to or from a memory location or an I/O port.

The Multibus interrupt lines consist of eight interrupt request lines and one interrupt acknowledge line. Interrupts are requested by activating one of the eight interrupt request lines. The interrupt acknowledge signal is generated by the bus master when an interrupt request is received. It effectively freezes the interrupt status and requests the placement of the interrupt vector address onto the data lines. There are six bus exchange lines that support two bus arbitration schemes on the Multibus system bus. A bus master gains control of the bus through manipulation of these signals. The bus request, bus priority, bus busy and bus clock signals provide for a slot dependent priority scheme to resolve bus master contention on the Multibus interface. Up to eight bus masters can be supported on the Multibus. Use of the common bus request signal line can save arbitration time by providing for a higher priority path to gain control of the system bus.

iLBX Structural Features

The iLBX uses a non-multiplexed 16-bit configuration capable of 8 and 16-bit transfers. Used in conjunction with the Multibus interface, the iLBX bus resides on the Multibus form factor P2 connector and supercedes the Multibus interface definitions for the P2 signals. The iLBX bus uses the standard 60-pin Multibus P2 connector and occupies 56 of the P2 connector pins with 16 data lines, 24 address lines plus control, command access, and parity signals. The four Multibus address extension lines on the Multibus iLBX P2 connector retain the standard Multibus interface definition.

iLBX Bus Elements

The iLBX bus supports three distinct device categories: 1) Primary Master, 2) Secondary Master, 3) Slave. These three device types may be combined to create several iLBX local buses ranging in size from a minimum of two to a maximum of five devices per iLBX bus. There is only one Primary Master in any given implementation of iLBX bus, and its presence is required along with the attachment of at least one Slave device. To provide alternate access over an iLBX bus, one optional Secondary master may be incorporated to create a "two-master" local bus subsystem. By limiting the iLBX bus to two Masters (a Primary and a Secondary), bus arbitration is reduced to a simple request and acknowledge process, with privileged use of the bus maintained by the Primary Master, and limited access granted to the Secondary Master when needed.

LAN Boards For The Multibus

A recent draft of the IEEE 802 states that the applications environment for the local area network is intended to be commercial and light industrial. Use of the local network in home or heavy industrial environments, while not precluded, is not considered. Specifically, the use of the local network for process control and other real-time, high reliability applications is also not considered. Taking into account the kind of environments that both Interlan and Proteon are selling into at the moment the statement would appear to be wrong.

At present, most R & D groups are creating a prototype and testing the technology before applying it. Many people are also using networks in the CAD/CAM area to link together expensive graphics processors. A third market, the one not discussed in the IEEE document is in process control. Although there has been some debate about Ethernet in the PC industry, some are using it as a channel between front-end processors and back-end devices like the PDP-11.

Earlier networks were stars, such as airline reservation systems that access central processing systems where communications were from user to mainframe. "That's getting to be passé, when you look at what's happening now," says Howard Salwen at Proteon. "Over at the lab for computer science they are experimenting with one VAX per user—you don't have a star network anymore but diffuse computer power spread over the floor."

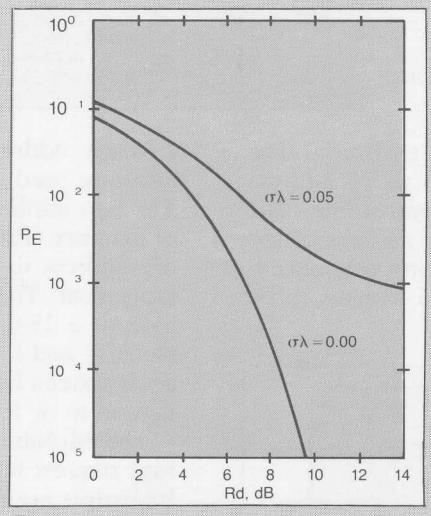

Figure 1: Average Probability of Error vs. Signal-to-Noise Ratio for Manchester Encoded Data.

Interlan and Proteon are offering different networking schemes to the marketplace both compatible with the Multibus. Interlan announced its first three major products, the NMIO Ethernet Protocol Module, NI 1010 Unibus and NI 2020 Q bus controllers in November 1981. The companies' fourth product was introduced in 1982. The NI310 Multibus Ethernet communications controller interfaces Multibus-based systems to the Ethernet local area network. It is specifically designed for OEMs implementing 68000, Z8000 and 8086-based systems. Interlan's second generation of Ethernet controllers, "the A-series," are claimed by the company to be the only group of products to support high station performances with DMA transfer to/from host memory.

In May, 1982 Interlan announced Etherway, an interconnection between DEC net and Ethernet. The NTIO Ethernet Transceiver featuring a non-intrusive cable top was introduced in September 1982. "Each transceiver top is a low capacitative load on the Ethernet—ours is 6pF at 250K, the lowest impedance on the marketplace," says Bob Olsen at Interlan. In a perfect world, one would like to pull as much energy as possible out of the cable, but one cannot because with Ethernet you have to share among the other users. This is why Interlan's transceiver aimed for such a low capacitive load.

Howard Salwen at Proteon was worried about the effects of transceiver reflections on the performance

The Primary Master executes the role of iLBX bus "supervisor" by controlling the general operation of the bus and managing Secondary Master accesses to the Slave memory resources.

The Secondary Master Device is an option providing alternate access to the Slave resources on the iLBX bus. Secondary Master Devices are typically DMA driven. This feature is provided for implementation flexibility when occasional DMA transfers in and out of iLBX memory resources can optimize the overall system performance. The Secondary Master essentially duplicates the Primary Master's data transfer capability, but must rely on the Primary Master to grant access. Once access is granted, the Secondary Master controls the bus, and drives all signal lines until the operation is complete and control is passed back to the Primary Master.

The Slave devices contain the memory resources used by the Primary Master and the optional Secondary Master. Each iLBX implementation can contain a maximum of four Slave devices. Using 64K RAM technology on four slave devices with ECC can provide for over 2 megabytes of "on board" high performance memory. With 256K RAM chips, each iLBX bus could contain slave devices with memory totalling 8 megabytes. As memory technology increases, the iLBX bus is designed to incorporate it in rapid fashion because it is capable of directly accessing a full 16 megabytes of memory on its high performance Slave devices.

of Manchester encoded systems, such as Ethernet, claiming that it was a cause of many missed Packets. He made his point in a recent article, and from that **Figure 1** shows the effect of synchronization error on probability of error performance in a Manchester encoded system. When the system is running at 10 MB/s, the estimate of bit synchronization must be less than 5ns RMS. Mr. Salwen states that such levels of performance are difficult to achieve in the burst mode. "I think it's important to know that there are problems with Ethernet—I think that software will bury these problems, and most people that use networks will not know that there are a lot of retransmissions going on," he says.

Out of the 500–1000 10 MB/s systems that are presently out in the field, Salwen claims to have greater than 50, 5/10% of the 10 Mbit/s market. His LAN, Pronet, is a token ring that operates at 10 MBit/sec. A maximum of 255 uses may be supported by each ring. **Figure 2** shows an extension of the star-shaped ring concept. Short runs from a wire center to various computers might be connected with twisted pairs. But long runs from wire center to wire center could be implemented with duplex fibre optic links. The attachment point for such fiber optic links at the wire centers includes a bypassing relay. If the fiber optic link connecting the two wire centers shown in **Figure 2** were to fail, the systems would automatically partition into two star-shaped rings with no interconnecting fiber optic link between them.

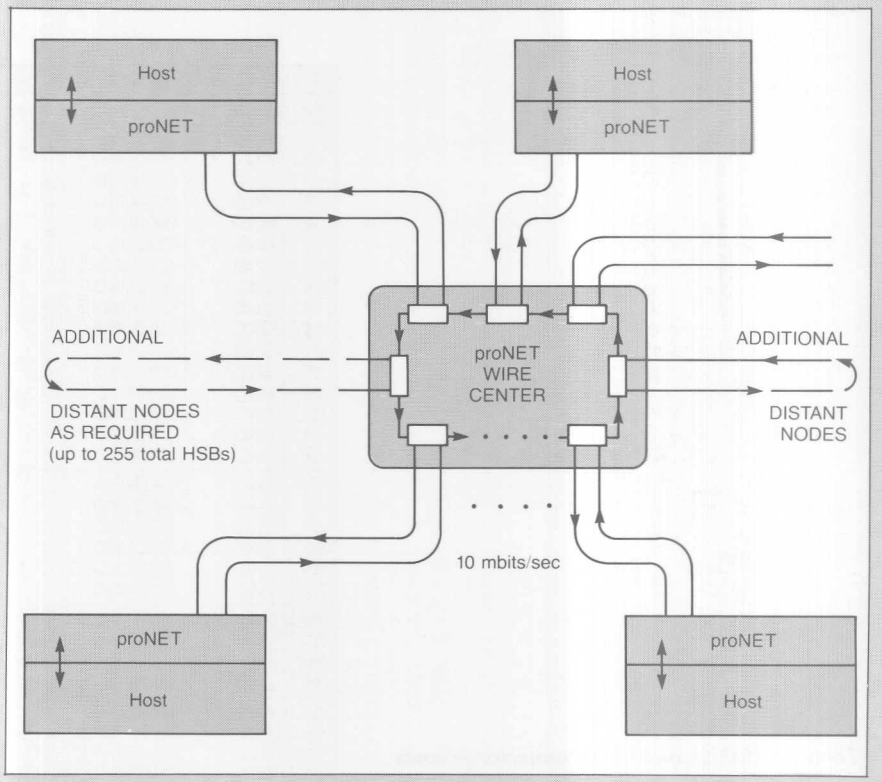

Figure 2: Basic Star-Shaped Ring.

iLBX Bus Interface/Signal Line Descriptions

The iLBX interface is divided into four functional classes of signal lines: address and data lines, control lines, command lines, and bus access lines. The 40 address and data lines defined by the iLBX Bus Specification consists of 16 data lines and 24 address lines. There are 16 bi-directional data lines exclusively used to handle 8-bit and 16-bit data transfers between the active bus master and the selected slave device. The iLBX bus uses these data lines for all data transfers, and are driven by tri-state drivers. The 24 address lines on the iLBX bus provide the ability to directly address 16 Mbytes of memory. These single-direction address lines are exclusively driven by the active bus master. The iLBX bus master uses them to select a specific slave device. Three control lines specify the type of data transfer between master and slave devices, while the three command lines initiate, control, and terminate the transfer. There are also three bus access lines used to transfer bus control between master devices.

iLBX Bus Pin Assignments

The iLBX uses the standard 60-pin Multibus P2 connector. The physical location of each pin asignment and its corresponding function is listed in **Table 2.** The four Multibus address extension lines (pins 55–58 on the P2 connector) retain the standard Multibus interface functions.

iLBX Bus Operation Protocol

The operation protocol for the iLBX bus is a straightforward set of procedures consisting of three basic operations: bus control access, write data to memory, read data from memory. These operations use asynchronous protocol with positive acknowledgement.

iLBX Bus Access

The iLBX is shared by at most two masters; one Primary Master and one Optional Secondary Master, each providing an alternate access path to iLBX bus memory resources. The mechanism for obtaining bus access is a simple request and acknowledgement process communicated between masters. Each master is a bus controller of similar capabilities, responsible for data transfer operations between devices, but the Primary Master has the added responsibility of controlling iLBX bus accesses. The Primary Master has default control of the iLBX bus. If the Secondary Master needs access to the bus, it must initiate a request and

LOWER ROW			UPPER ROW		
Pin	Mnemonic	Signal Name	Pin	Mnemonic	Signal Name
1	GND	GROUND	2	AD0/	ADDRESS DATA LINE 0
3	GND	GROUND	4	AD1/	ADDRESS DATA LINE 1
5	GND	GROUND	6	AD2/	ADDRESS DATA LINE 2
7	GND	GROUND	8	AD3/	ADDRESS DATA LINE 3
9	GND	GROUND	10	AD4/	ADDRESS DATA LINE 4
11	GND	GROUND	12	AD5/	ADDRESS DATA LINE 5
13	GND	GROUND	14	AD6/	ADDRESS DATA LINE 6
15	GND	GROUND	16	AD7/	ADDRESS DATA LINE 7
17	GND	GROUND	18	AD8/	ADDRESS DATA LINE 8
19	GND	GROUND	20	AD9/	ADDRESS DATA LINE 9
21	GND	GROUND	22	ADA/	ADDRESS DATA LINE 10
23	GND	GROUND	24	ADB/	ADDRESS DATA LINE 11
25	GND	GROUND	26	ADC/	ADDRESS DATA LINE 12
27	GND	GROUND	28	ADD/	ADDRESS DATA LINE 13
29	GND	GROUND	30	ADE/	ADDRESS DATA LINE 14
31	GND	GROUND	32	ADF/	ADDRESS DATA LINE 15
33	GND	GROUND	34	RESET/	RESET
35	GND	GROUND	36	AACC	ADDRESS MODE ACCEPT
37	GND	GROUND	38	SRQ/	SERVICE REQUEST
39	GND	GROUND	40	STO/	SUPERVISOR TAKE OVER
41	GND	GROUND	42	DACC/	DATA MODE ACCEPT
43	GND	GROUND	44	SA/	SUPERVISOR ACTIVE
45	PB*/	PARITY BIT (INV.)	46	PB/	PARITY BIT
47	R/W/	READ NOT WRITE (INV.)	48	R/W	READ NOT WRITE
49	A/D/	ADDRESS NOT DATA (INV.)	50	A/D	ADDRESS NOT DATA
51	DYDY*/	DATA READY (INV.)	52	DRDY/	DATA READY
53	RES	RESERVED	54	RES	RESERVED
55	RES	RESERVED	56	RES	RESERVED
57	RES	RESERVED	58	RES	RESERVED
59	RES	RESERVED	60	RES	RESERVED

Table 3: Multichannel bus connector pinouts.

wait for acknowledgement from the Primary Master. The choice of when to surrender control of the bus rests with the Primary Master, but if no data transfer is in progress, the Primary Master normally relinquishes control immediately to the Secondary Master.

Data Transfer Operation

The iLBX bus supports two types of data transfer operations: write data to memory and read data from memory. These data transfer operations facilitate the passing of information between the active bus master and the selected slave device. The operation of these two transfer types is very similar; the only differences being the direction of the data transfer and the device driving the data lines. For either type of data transfer, the active bus master first initiates the transfer operation by placing the memory address on the address lines (AB23-AB0) and a control configuration on the control lines to select the slave device. Once the slave device is selected, the type of data transfer becomes the key factor. With the write operation, the active master maintains control of the data lines and provides valid data within the specified time. Upon accepting a data element, the slave sends a receipt acknowledgement signal to the master which completes the data transfer operation.

With the read operation, the slave device drives the data lines and places valid data on the data lines before sampling by the active master. The slave acknowledges the master to signal the end of the data transfer, and the master completes the operation.

The iLBX Bus Specification includes provisions for both optimized and non-optimized data transfers. Optimized operation uses pipelining and signal overlapping techniques to manage the data transfer timing relationships between the active bus master and the selected slave. The use of signal overlapping requires that every device attached to the iLBX bus provide a means for varying the timing of the slave request and acknowledge signals. The non-optimized operation uses fixed signal sequences, instead of signal overlapping, to assure a valid data transfer, and a device does not need a variable request or acknowledge to read data-valid timing on the iLBX bus.

Mechanical Implementation

Because the iLBX bus uses the P2 connector of the Multibus form factor, the iLBX bus "shares" a Multibus chassis with the Multibus backplane system bus in the system design. The iLBX mechanical specifications are synonymous with the Multibus specifications for board-to-board spacing, board thickness, component lead length, and component height above the board.

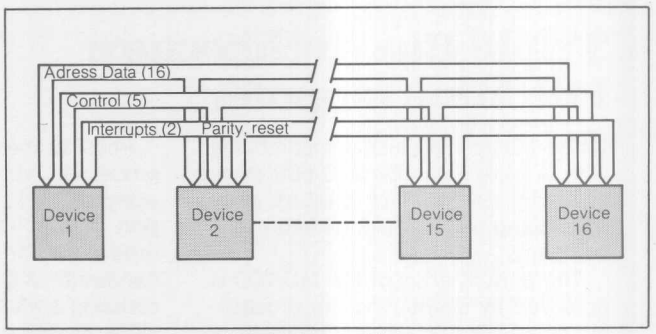

Figure 7: Structure of Multichannel bus.

The ILBX bus interconnection can either use flexible ribbon cable or a rigid backplane. The iLBX bus interconnect maximum length is limited to 10 cm (approximately 4 inches); that is sufficient to span 5 card slots across two connected chassis.

Multichannel I/O Bus

The Multichannel bus is a multiplexed, asynchronous, block transfer, 16-bit I/O bus designed to handle 8-bit and 16-bit transfers between peripherals and single board computers. Its structure (**Figure 7**) consists of 16 address/data lines, 6 control lines, 2 interrupt lines, plus parity and reset. Signals are implemented as either a 60 conductor flat ribbon cable or a twisted-pair cable spanning a distance of up to 15 meters. A 30/60-pin 3M connector is recommended for device connection to the Multichannel bus. The male connectors are installed on each Multichannel device and the female connectors are mounted on the cable. To insure system integrity, the Multichannel cable is terminated at both ends.

Multichannel Bus Elements

Three device types—the Basic device, the bus Controller device, and the bus Supervisor device—each provide a different level of capability. The Basic Talker/Listener device has lowest capability, responding only to data transfer requests issued by a Supervisor or Controller. The bus Controller device has higher capability than a Basic Talker/Listener on the bus. It can respond to data transfer requests, control data transfers, and can program other Multichannel devices under direction from a bus Supervisor. Operating at the highest capability is the bus Supervisor device. It

Triple Bus Architecture For Minicomputer

Using Multiple processors combined with a triple bus, the MC-500 minicomputer is the first product from Masscomp, a year-old Littleton, MA company.

The performance of the MC-500 is achieved by distributing the processing capabilities to several processors within the system. The MC-500 incorporates twin 68000's, a third 68000 in a separate display processor subsystem, and a bit-slice machine in a data acquisition system. The minicomputer can be further enhanced with an integrated floating point array processor.

Physical memory is attached to the processor via Masscomp's proprietary MC–500 bus, designed to support 4-byte transfers using a block mode protocol. Memory system bandwidth is 8 Mbytes/sec. The CPU contains a Multibus adapter that provides an I/O map that translates the address of any DMA memory reference to an appropriate physical page number. The map contains translations for 1024 pages. A special 32-bit transfer mode enables the Multibus to offer a 6 Mbyte/sec. bandwidth.

Masscomp chose the Multibus as the MC–500 system peripheral bus because of the I/O options that are available. Masscomp supports multiple disk and multiple ASCII terminal configurations, network hardware and software, printers and plotters as well as the Masscomp Data Acquisition and Control Processor plus up to four Masscomp Independent Graphics Processors.

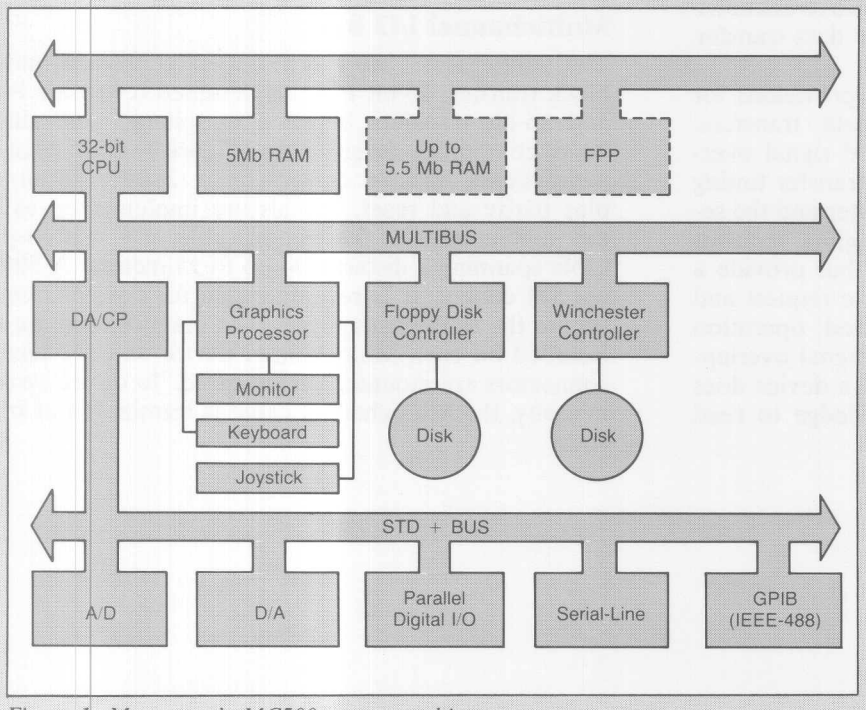

Figure 1: Masscomp's MC500 system architecture.

System Structure

The operating system for the MC–500 is based on Unix System III, enhanced by Masscomp to provide real-time services. Virtual memory support and other features were added using the UC Berkeley Unix extensions.

In addition, Masscomp's UNIX provides memory locked processes to prevent real-time processes from being swapped out. This ensures a predictable system response to external interrupts as well as ensuring that graphics display lists and data buffers remain in physical memory during critical operations.

Data Acquisition And Control

The DA/CP is based on a high-performance bipolar bit-slice processor that executes each instruction in 125ns. In addition to this processor, the DA/CP includes 1024 locations of 24-bit temporary data storage, two 16-element FIFOs, and one 64-element FIFO. These components enable devices on the STD bus to input or output data to data buffers within an application program.

STD+ Bus

The STD+ Bus was developed by Masscomp to enhance the industry

provides major control and management of the Multichannel bus. The bus Supervisor resolves and grants Multichannel bus priority, monitors bus status, handles interrupts, and controls the reset line, in addition to performing all bus Controller functions.

Multichannel bus devices are functionally flexible, creating overlaps between types of bus functions and types of bus devices performing those functions. These devices perform functions in various states of operation: master, slave, talker, listener. When a device is controlling the command/action lines, it is in the master state, and both the bus Supervisor and the bus Controller can operate in this state, although not simultaneously. The slave state indicates a device that can monitor the command/action lines. Only Controllers and Basic Talker/Listeners operate as slaves. All three device types can operate in the talker state or the listener state, but not all at the same time. A Talker is any device selected by the bus master which is writing data to the bus. A Listener is any selected device which is reading data from the bus.

Multichannel Bus Interface/Signal Line Descriptions

The Multichannel bus signal lines are grouped into five basic classes based on the functions they perform: address/data, control, interrupt, parity, and reset. The 16 address/data lines are multiplexed by a control line to act either as 16 unidirectional address lines or 16 bidirectional data lines. When used as address lines, they transmit the device address to all devices attached to the Multichannel bus. When used as bidirectional data lines, they transmit and receive data to or from Multichannel devices. The six control lines determine the overall operation of the bus from specifying the type of data transfer to providing the handshake for data transfers between Multichannel devices. Two in-

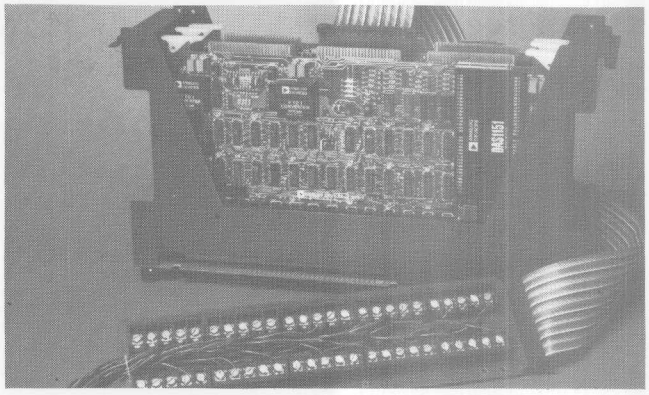

Analog Devices' RTI-711, RTI-724 and RTI-732 bring data acquisition and control capability to Multibus compatible microcomputers.

standard STD bus. The STD+ Bus consists of two side-by-side 9-slot STD Buses that share address lines but have separate data paths and read/write control lines. Most Z80 compatible STD boards will plug into the STD+ Bus. The data transfer rate of these devices is limited to the STD Bus maximum of 1 million bytes/second.

The STD+ interfaces from Masscomp are physically twice as high as conventional STD boards, and therefore, plug into two slots. Because the data are transferred from Masscomp interfaces using two seperate byte-wide paths, the STD+ Bus provides a data rate of up to 2 million bytes/second.

Analog And Digital Interfaces

Masscomp offers a wide range of analog and digital interfaces for the STD+ Bus. The analog interfaces include a 40kHz and a 1 MHz A/D converter, each with a resolution of 12-bits. Both interfaces provide 16 single-ended or 8 different channels, and both may be expanded to provide 64 single-ended or 32 differential channels. The 1 MHz interface includes a programmable gain capability and is available with an 8-channel Simultaneous Sample-and-Hold option. As a complement to these interfaces, Masscomp offers 4 and 8 channel D/A converters which can output data at a rate of 500kHz per channel.

Masscomp digital interfaces include a 16-line parallel interface, a dual-port RS–232C serial-line interface, and a GPIB(IEEE–488 Bus) interface. These digital interfaces are able to support a wide range of laboratory instrumentation.

Pin[1]	Menmonic	Description	Pin[1]	Mnemonic	Description
43	MD8	MDATA Bit 8	44	MD9	MDATA Bit 9
41	MDA	MDATA Bit A	42	MDB	MDATA Bit F
39	MDC	MDATA Bit C	40	MDD	MDATA Bit D
37	MDE	MDATA Bit E	38	MDF	MDATA Bit F
35	GND	Signal Gnd	36	+5V	+5 Volts
33	MD0	MDATA Bit 0	34	MDRQT	M DMA Request
31	MD1	MDATA Bit 1	32	MDACK/	M DMA Acknowledge
29	MD2	MDATA Bit 2	30	OPT0	Option 0
27	MD3	MDATA Bit 3	28	OPT1	Option 1
25	MD4	MDATA Bit 4	26	TMDA	Terminate DMA
23	MD5	MDATA Bit 5	24		Reserved
21	MD6	MDATA Bit 6	22	MCS0/	M Chip Select 0
19	MD7	MDATA Bit 7	20	MCS1/	M Chip Select 1
17	GND	Signal Gnd	18	+5V	+5 Volts
15	IORD/	I/O Read Cmd	16	MWAIT/	M Wait
13	IOWRT/	I/O Write Cmd	14	MINTR0	M Interrupt 0
11	MA0	M Address 0	12	MINTR1	M Interrupt 1
9	MA1	M Address 1	10		Reserved
7	MA2	M Address 2	8	MPST/	iSBX Multimodule Board Present
5	RESET	Reset	6	MCLK	M Clock
3	GND	Signal Gnd	4	+5V	+5 Volts
1	+12V	+12 Volts	2	−12V	−12 Volts

Notes:
1. Pins 37–44 are used only on 8/16-bit systems
2. All undefined pins are reserved for future use.

Table 4: iSBX signal/pin assignments.

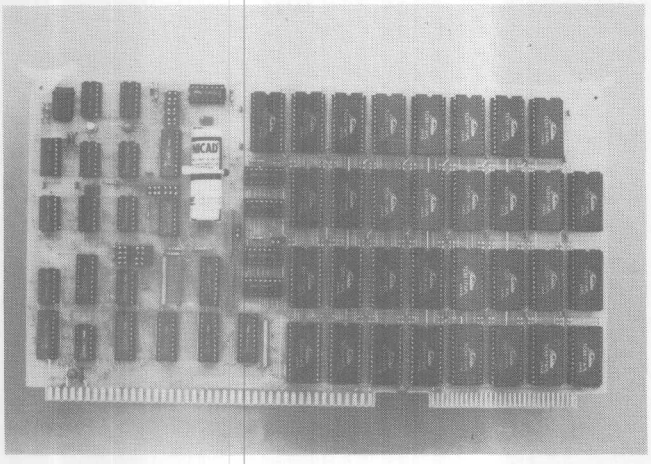

Diversified Technologies' 64k CMOS Battery Backup RAM Board is aimed at applications in instrumentation, process control and remote or portable instruments.

terrupt lines are supplied to initiate and terminate data transfers, and to indicate device failures, memory failures, or parity errors. A parity line and a reset line provide support for a parity option and system reset capability whenever required.

Multichannel Bus Pin Assignments

For proper Multichannel implementation, a 60 conductor (twisted pair or flat) cable using a 30/60 pin 3M connector, is used for device connection to the bus. The Multichannel bus connector signal pin assignments are listed in **Table 3**. Cable termination is implemented at both cable ends to insure proper system integrity over a 15-meter cable.

iSBX Bus Elements

The iSBX Multimodule system is made up of two basic elements: base boards and iSBX Multimodule boards. In an iSBX system, the role of the base board is simple. It decodes I/O addresses and generates the chip selects for the iSBX Multimodule boards.

The iSBX bus supports two classes of base boards, those with direct memory access (DMA) support and those without. Base boards with DMA support have DMA controllers that work in conjunction with an iSBX Multimodule board (with DMA capability) to perform direct I/O to memory or memory to I/O operations. Base boards without DMA support use a subset of the iSBX bus and simply do not use the DMA feature of the iSBX Multimodule board.

The iSBX Multimodule boards are small, specialized, I/O mapped boards which plug into base boards. The iSBX boards connect to the iSBX bus connector and convert iSBX bus signals to a defined I/O interface.

iSBX Bus Interface/Signal Line Descriptions

The iSBX bus interface can be grouped into six functional classes: control lines, address and chip select lines, data lines, interrupt lines, option lines, and power lines. The iSBX provides nine control lines that define the communications protocol between base board and iSBX Multimodule boards. These control lines are used to manage the general operation of the bus by specifying the type of transfer, the coordination of the transfer between devices. The five address and chip select signal lines are used in conjunction with the command lines to establish the I/O port address being accessed, effectively selecting the proper iSBX Multimodule. The data lines on the iSBX bus can number 8 or 16, and are used to transmit or receive information to or from the iSBX Multimodule ports. Two interrupt lines are provided to make interrupt requests possible from the iSBX board to the base board. Two option lines are reserved on the bus for unique user requirements, while several power lines provide +5 and ±12 volts to the iSBX boards.

iSBX Bus Pin Assignments

The iSBX bus uses widely available connectors that are available in 18/36 pin for 8-bit devices and 22/44 pin for 16-bit devices. The male iSBX connector is attached to the iSBX Multimodule board and the female iSBX connector is attached to the base board. A unique scheme allows the 16-bit female connector to support 8 or 16-bit male Multimodule boards. **Table 4** lists the signal/pin assignments for the bus.

Summary

All of these structures form a consistent architecture for high performance, modular systems.

It is unique in how much of the system interfaces are standardized, ensuring compatibility with future technology. All eyes now rest on Intel to see how the Multibus II evolution will affect the integrator of future 32–bit systems, and just how compatible it will be with Multibus I.

For More Information

A complete up-to-date buyers guide of Multibus manufacturers is available from Ironoak Co., 3239 Caminito Ameca, La Jolla, CA 92037. Tel: 714-450-0191. □

* *Multibus is a registered trademark of Intel Corp.*

ARTICLE
REPRINT

AR-267

April 1983

Choosing a Bus for Control

John Beaston
OEM Modules Operation
MULTIBUS® Architect

"Reprinted with permission from Instruments & Control Systems, March 1983 issue."

230727-001

Choosing a bus for control

Which system bus is best for your application? Both business and technical factors should be considered when making a selection.

John Beaston
Multibus Architect
Intel Corporation
Hillsboro, OR 97123

Last month's installment of our "Bus boards for control" series provided a broad overview of bus-based systems, including a look at their features, advantages and disadvantages. This second installment describes the business and technical factors that must be considered when selecting a bus for your application from the dozens available.

• • • • •

The importance of selecting the best bus for your application can't be overemphasized. This is true whether you're working for an OEM or are building a one-of-a-kind product for internal use. The choice you make can impact such important areas as system performance, the engineering time needed for system configuration (which affects time-to-market for OEM's), and the ability to enhance the system in the future.

To help you with this critical decision, this article looks at both the business issues and technical aspects that should be considered when choosing a bus. It uses six of the more popular buses available to make example comparisons. These include Multibus, VME, Versabus, STD, S-100, and the Q-bus.

Business factors

Because many of the buses have similar technical capabilities, business factors often sway the final decision in bus selection. The most important of these are described in the material that follows. (Table 1 compares business factors for the six buses mentioned earlier.)

Public specification

Choosing a bus with publicly available specifications can provide some measure of compatibility among the board products supplied by different vendors. Unless all board designers use the same spec, compatibility is a hit-or-miss proposition.

Controlling body

Having a public spec, however, does not guarantee compatibility. Public specs can change with time. There-

Table 1: Business factors for various buses

	Multibus	VME	Versabus	STD	S-100	Q-bus
Public spec	Yes	Yes	Yes	Yes	Yes	Yes
Controlling body	IEEE 796	Motorola, Signetics, Mostek	Motorola IEEE P970*	IEEE P961	IEEE 696	DEC
Multiple vendors	150	21	10	75	100	10
Second sources	Yes					Yes
VLSI support	Yes	Yes	Yes			Yes
Processors	8080, 8085, 8086, 80186, 80188, 80286, Z80, Z8000, 68000, 16032, 6800, 6100, NCS800, 6802, 6808	68000, 80186, 16032	68000	8085, 8088, 6800, Z80	8080, 8085, 8088, 80186, 80286, Z80 6800, 68000, 16032	LSI-11

*Pending

fore, to prevent incompatibility and give bus users confidence, any changes made should receive an impartial review by a single controlling body.

A single point of control is critical. For example, until it was adopted by the IEEE, the S-100 bus had no control at all, let alone from a single point. Stores of incompatible S-100 boards are well-known. Now, with IEEE control and standardization, compatibility should be assured, at least for new products.

Multiple vendors

No one manufacturer can serve all the needs of the bus board marketplace. And the best bus specification in the world doesn't mean a thing if the type of board you need isn't available. In general, the number of products available on a given bus is proportional to the number of vendors supporting it. Thus, selecting a bus with strong multiple vendor support can be to your advantage.

Second sources

Some bus manufacturers are second sourcing the more popular boards. This helps assure board availability, and can also mean competitive pricing.

Processor independence

One issue that can affect future system enhancements is processor independence. Some buses are specific for a particular microprocessor family or type of processor. This can lock you into a processor family or type even though technology advancements offer better solutions. A bus which supports many processors lets you choose the best processor for your application without having to throw away your investment in the bus.

VLSI support

The final business issue is VLSI (very large scale integration) support for the bus. Some manufacturers have reduced the interface to the buses they support to a handful of VLSI devices. Having these available can greatly simplify the job of building a board, and can improve its performance and functionality by freeing up more board real estate for useful functions. It can also make off-the-shelf boards less expensive.

Technical factors

Let's turn now to a look at the technical factors in bus selection. Areas to be discussed include: address range, data width, bandwidth, interrupts, multiple bus masters, arbitration, mutual exclusion, board size, connector type, and I/O module bus. (Table 2 compares these factors for the six buses mentioned earlier.)

Address range

A system's memory requirements tend to go up as its capabilities and performance requirements increase. Therefore, you should choose a bus which supports an address space at least as large as the highest level processor you expect to use on it. This is generally a good idea even though at first, you might only be using a small portion of the address space. It will allow for later expansion.

On buses which specify expanded or optional address ranges, be very careful when selecting board products. Be sure the product decodes all of the address range you need. For example, if you use a 24-bit address range and buy a memory board which decodes only 20 bits, you'll find its address space is duplicated every 1 Mbyte on your 16 Mbyte space—which can be very inconvenient. Some buses, like VME and Versabus, have a way around this.

Data width

A processor presents a certain number of external data pins to the outside world and the bus you select should match this. While it is possible to put a 16-bit processor on a 8-bit wide bus and have the bus do two transfers each time the processor requests 16 bits, it's slow, clumsy, and complicated.

Bandwidth

Table 2 gives the maximum theoretical bandwidth for each bus at various data widths. Note that, with two exceptions, the bandwidths for any given data width are within about 20% of each other. This is because all the buses dis-

Table 2: Technical factors for various buses

	Multibus	VME	Versabus	STD	S-100	Q-bus
Address width	24	16, 24, 32	16, 24, 32	16	16 standard 24 expanded	16 standard 22 expanded
Data width	8 | 16	8 | 16 | 32	8 | 16 | 32	8	8 | 16	8 | 16
Bandwidth (Mbytes/s)	5 | 10	6 | 12 | 24	6 | 12 | 24	1	6 | 12	0.8 | 1.6
Interrupt lines	8	7	7	2	10	4
Interrupt ack	Polled	Daisy-chain	Daisy-chain	Daisy-chain	Polled	Daisy-chain
Mutual exclusion	Bus lock	RMW+bus lock	RMW+bus lock	None	Bus lock	None
Arbitration	Serial or parallel	Serial or parallel with daisy-chain	Parallel with daisy-chain	Serial	Parallel	Serial
MECHANICAL						
Form factor	6.75x12	6.3x9.2	9.25x14.5	4.5x6.5	5.1x10	5.25x8.9
Area (in.²)	81	58	134	29	51	47
Connector size	86/60	96/96	140/120	56	100	36/36
Connector type	Edge	DIN	Edge	Edge	Edge	Edge

cussed in this article use an asynchronous bus protocol with roughly the same timing. The only exceptions are the STD and Q-bus, which have relaxed timing specs.

The bandwidth figures given in Table 2 ignore bus arbitration delays and assume zero memory access times. But memory access times constitute over half the time of a typical bus transfer. They reduce the available bandwidth by more than 50%, and make the comparison come out even closer. Only the STD and Q-bus are restrained by their bus timing. The others are at the mercy of their memory boards. If you're looking for maximum performance, pick the bus that enables you to buy the fastest memory boards.

Interrupts

Systems are usually divided along functional lines. For example, a control system might be partitioned into analog I/O, digital I/O, and operator console functional modules. These modules could all be physically located on the same board and controlled by a single processor, or each module might be on a different board with individual control processors. However they're arranged, the modules must communicate.

The traditional method is with interrupts. Ideally, each module would have its own interrupt line, and each line would have a different priority, as shown in Fig. 1. Since often there are too many modules for this to be practical, bus designs are the result of a compromise. They allow the modules to share the interrupt lines, leaving it up to the interrupted device to figure out the source of the interrupt.

There are two main ways of doing this: software polling and daisy-chained acknowledge.

In software polling (Fig. 2), the interrupted module polls, or asks, each of the modules on the applicable interrupt request line. The module which generated the interrupt replies when polled.

Daisy-chained acknowledge (Fig. 3) is a hardware method. An interrupt acknowledge signal is passed down a line daisy-chained between all modules on the shared interrupt line. Modules not signalling the interrupt let the acknowledge signal pass on down the daisy-chain. Once the signal arrives at the interrupting module, the module captures it and places an identifying code on the data bus. The interupted module then reads this code to figure out the source.

It's hard to say whether one method is better than the other. Polling, being software based, is very flexible and easy to debug, but tends to be slower than a hardware method. Daisy-chaining is fast but makes the priorities of modules sharing an interrupt line depend on their positions.

If there are fewer modules than the number of interrupt lines, neither method offers an advantage.

Multiple bus masters

If there is one characteristic which tells if a bus is suitable for mid-to-high end systems, it is its support of multiple bus masters. Such support greatly enhances the architectural flexibility of the bus.

Simple buses like the STD and Q-buses support only one bus master in addition to the main processor. This usually must be a direct memory access (DMA) device. Typically, both masters can't operate at the same time. If the DMA device is active, the main processor suspends operation temporarily even though it might be using only on-board private resources. Once the DMA device is done with the bus, the main processor resumes operation. This nonconcurrency generally restricts these buses to low-to-mid range performance systems.

The remaining buses support multiple bus masters in the true sense of the term. On these buses, it's possible to have many processors sharing the bus. Usually the processors operate with their on-board resources while sharing a common global pool of memory and/or I/O.

Multiple-master buses give you a lot of flexibility in system architecture. A simple system might use a single main processor board with a dumb analog I/O board. Or you could increase the system's throughput by replacing the dumb I/O board with an intelligent board containing its

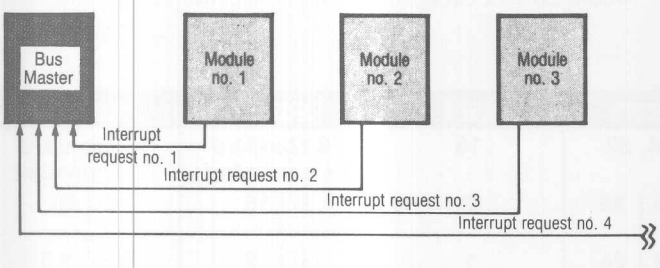

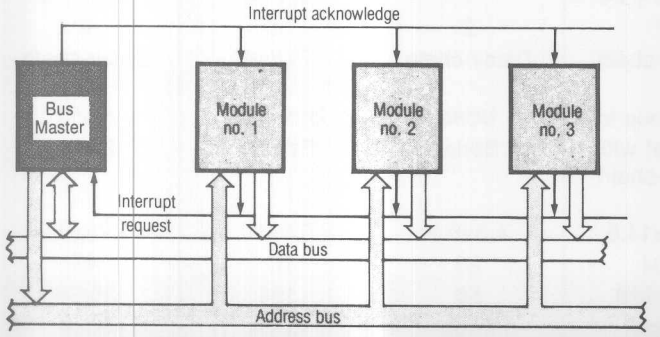

Fig. 1 (left, top): If there are more interrupt request lines than interrupting modules, the sources of the interrupts are obvious.

Fig. 2 (left, bottom): When modules must share interrupt request lines, there must be some way to find the source of an interrupt. In software polling, the bus master asks each module in turn if it is the source.

Fig. 3 (below): In daisy-chained acknowledge, the module which caused the interrupt captures the acknowledge signal to indicate that it is the source. Priority then depends upon position on the bus.

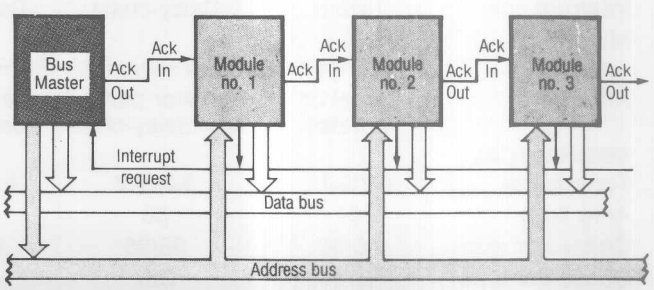

own dedicated processor. This intelligent board would operate in parallel with the main processor. The support of multiple masters lets a bus span a tremendous cost/performance range.

Arbitration

All of the buses mentioned only allow one bus master at any one time. Since several masters might want the bus at the same time, arbitration is needed to resolve the conflict and figure out which one gets the bus and which one has to wait.

This can be done with either serial or parallel methods. In parallel arbitration, each bus master (or group of masters in those buses with daisy-chained arbitration) generates a bus request, which goes to a central arbitration module. This module then resolves any simultaneous requests based upon a priority algorithm and returns a grant to use the bus to the winning master. In a purely parallel scheme, there are as many separate request and grant lines as there are masters. However, buses like VME and Versabus allow bus masters to share common request lines. The grant for each request line is then daisy-chained to determine which of the masters gets the bus.

All parallel arbitration schemes require external arbitration logic. Sometimes this logic is placed on the backplane, other times it's included on a processor board, and still other times a separate board is used. How it's done is an implementation detail that you should be aware of when you choose board products for your selected bus.

To save the cost of parallel arbitration for simpler systems, the Multibus uses a serial arbitration scheme. This requires no external logic but can only support three bus masters, which is often enough. VME can also do serial arbitration, with up to 20 masters.

Mutual exclusion

Systems which share common resources, such as a shared memory pool or a peripheral, need some way to prevent conflict. One popular method is to set a bit in memory called a semaphore to indicate that the resource is in use. A processor wanting to use the resource will read the semaphore; if it is not set, the processor will set it and use the resource.

But there can be a problem. After the bit is read, and before it can be set, another processor can read it, find it cleared, and try to set it. To avoid this problem, buses may use mutual exclusion.

The various buses provide mutual exclusion in different ways. VME and Versabus use read-modify-write (RMW) bus cycles. These are read and write cycles to the same memory location, which are treated as one bus transaction. No arbitration is allowed between the cycles so it's not possible for another processor to sneak on the bus in between read and write operations. RMW cycles are used for protecting semaphores. The semaphores, in turn, protect other shared data structures, such as pointers.

The other mutual exclusion mechanism is called bus lock. Here, when a processor wants exclusive use of a global bus resource it stops arbitration until the access is complete. The bus lock can cover more than just a read and write cycle. It could apply over many bus cycles if so desired.

It is also useful with dual-ported memory. It can inhibit access from any port of a multi-ported resource so that even though a local processor is using its local bus, it cannot violate mutual exclusion.

Board size

Perhaps the most important factor in determining the type of systems a bus supports is its form. Smaller boards are very modular. They allow you to get exactly what you need with few excess or unused functions. On the other hand, they may have lower performance. Since each function usually has its own board, and its own bus interface, more real estate goes to interfacing to the bus. Larger boards can have more functions per board, so the percentage of area devoted to the bus is lower. This leaves more room for extra functions.

The problem with larger boards is finding ones with the right mix of functions. Unused functions waste real estate, power, and money. Supporters of the larger boards (Multibus and Versabus) have recognized this problem and have developed I/O module bus structures as extensions to the primary system bus. We'll discuss these extensions shortly.

The other tradeoff on board size is performance. Many of the boards use a local bus. This is an on-board bus which the processor can use without disturbing the main system bus. These local buses tend to have very high performance, since the distances are short and arbitration is usually not required. It's good to maximize the time a processor spends on its local bus since it raises its overall performance. No matter how fast the system bus may be, the local bus is always faster.

In general, smaller boards have less room for local bus resources. This is particularly true for memory. Once the processor exceeds the amount of local memory, it must use the system bus and suffer the delays of arbitration and multiple logic gates. Larger boards have larger local memory capacity and shouldn't have to use the system bus as often; hence, their system performance can be higher. Multibus has a local bus extension (iLBX) available. This extends a board's local bus to other boards without going through the system bus.

Connector type

All of the buses in this comparison except VME use single-piece edge connectors. VME uses the European standard DIN pin-and-socket connector. Edge connectors have been around a long time and have acceptable connection reliability performance, but some people feel that pin-and-socket connectors are better in high-vibration environments. They are also gas-tight while edge connectors are not.

Another motivation for using DIN connectors is space efficiency. For connections larger than 100 pins, DIN connectors take up less board area than edge connectors. The small VME boards with fully demultiplexed 32-bit address and 32-bit data buses need DIN connectors to support all those pins.

Table 3: Comparison of I/O module buses

I/O Module Bus	iSBX	I/O Channel
Companion system bus	Multibus	Versabus VME
Public spec	Yes	Yes
Controlling body	IEEE P959	Motorola
Multiple vendors	10	2
Address width	5	12
Data width	8 | 16	8
Bandwidth (mByte/sec)	5 | 10	2
DMA support	Yes	
I/O connection type	Daughter board	Separate cable or backplane

The other major differences are cost and current capacity. Pin-and-socket connectors are more expensive than edge connectors, and they handle less current than edge connector fingers (1to 1.5 A per pin connection vs 1 to 5 A per edge connector finger). This means that buses using pin-and-socket connectors need to use more pins to supply power.

I/O module bus

As mentioned earlier, the supporters of the larger boards have developed specialized I/O buses to improve their modularity. Multibus uses the iSBX I/O Expansion Bus and VME/Versabus have the I/O Channel. Table 3 gives a comparison of these buses.

While the goals are similar, the methods are different. The I/O Channel uses a separate ribbon cable or backplane, and can extend up to twelve feet from the host system. Up to 16 slave devices can be connected, transferring data over an 8-bit wide path.

The iSBX is a small daughter board which plugs directly into a host Multibus board, and each daughter board handles one kind of I/O. Both 8 and 16-bit data transfers are possible.

A final word

With this information you should be able to make a more informed choice of a bus for your application. Just remember that there are a lot of buses available, not just the six discussed here. And there may be more than one which will suit a particular application. ■

The author, John Beaston, will be available to answer any questions you may have about this article. Mr. Beaston can be reached at (503) 681-8080 during normal business hours.

News

Extension steps up Multibus data transfers for 8- and 16-bit words

A local bus extension raises the speed of transfers over the Multibus to 9.5 Mbytes/s for 8-bit words and 19 Mbytes/s for 16-bit data.

The delays usually incurred when a processor board taps off-board resources have been surmounted with a local bus extension to the IEEE-approved Multibus. With the iLBX extension, Intel Corp. (Hillsboro, Ore.) has successfully boosted Multibus data rates to 9.5 Mbytes/s for 8-bit transfers and 19 Mbytes/s for 16-bit transfers. Furthermore, by using the local extension, the main processor gains a dedicated 16-bit bus, over which it can make memory transfers, and can also address up to 16 Mbytes of what appears to be on-board memory.

The new set of bus-control and data lines resides on the previously unused 60-pin P_2 edge connector adjacent to the standard Multibus connector. Of those pins, 40 are dedicated to a 24-line address bus and a 16-line data bus and others to various control, command, and access lines. Part of the original Multibus extension, which also handles 16 bits, uses four pins that were allocated on the P_2 connector.

Building the extension

The iLBX supports three types of boards: a primary master, a secondary master, and a slave. Only one primary master can be used in a system, and it should be complemented with a slave. The primary master supervises the local bus, controlling its general operation and managing the secondary master's access to memory resources. Slave boards are usually memories that contain circuitry for detecting and correcting data errors (see the figure). The secondary master, which is typically controlled by DMA needs, can be added for alternative access to the local bus. At any one time, the iLBX can work with five boards.

The bus protocol covers three basic operations: controlling access to the bus, writing data to memory, and reading data from memory. All three operations take place asynchronously and have a positive acknowledgment. The bus is also capable of optimized and nonoptimized data transfers. Optimized operations use pipelining and signal overlapping techniques to manage timing relationships between the active master board and the selected slave. Nonoptimized operations employ fixed, rather than overlapping, signal sequences to ensure valid transfers.

Dave Bursky

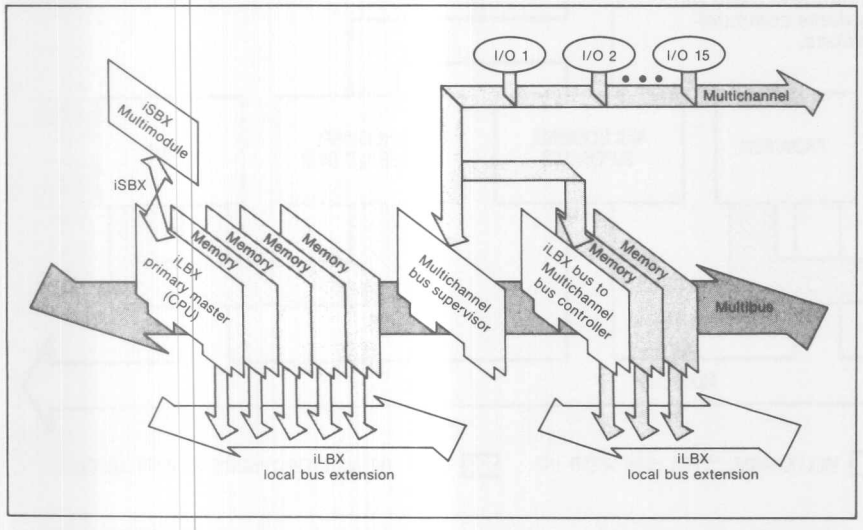

The iLBX local bus extension to the Multibus adds up to a 19-Mbyte/s dedicated memory bus, which appears as an on-board resource to the master CPU. This shortens the memory-transfer delay from the master to one of the memory boards.

230727-001

Multibus architecture extends with local bus for high-speed memory transfers

286-based SBC and DRAM cards implement new scheme

In our October update on single-board computers, we reported on efforts of several vendors to extend the capability of the Multibus architecture to take advantage of newer 16-bit processors and denser, faster memory chips. Now Multibus originator Intel has extended its architecture with a high-speed local bus for direct processor-to-memory data transfers. The iLBX bus uses the existing P2 connector to allow 9.5 Mbyte/s 8-bit transfers and 19 Mbyte/s transfers of 16-bit data. A group of 19 vendors has already endorsed the extension, including Microbar Systems (who implemented a similar approach last year) and IEEE standardization activity has already begun.

The iLBX is needed because even with 64K RAMs, fully functional single-board computers can only accommodate about 256 Kbytes of on-board memory for high-speed access. Given the wide address range and high-speed execution of second generation 16-bit devices like the 286 and the 68010, this is inadequate. The local bus makes up to four expansion memory boards with up to 16 Mbytes total appear as local memory to the CPU. Maximum bus length is 10 cm.

With the introduction of iLBX, Intel has now neatly partitioned interboard communications processes three ways and significantly extended the mature bus' usefulness and lifetime. Last year saw the introduction of the 8 Mbyte/s Multichannel I/O bus, a 16-bit parallel bus designed specifically to link up to 15 direct memory access devices (such as disk drives, graphics subsystems, etc.) to a Multibus system at distances up to 15 meters.

Using the Multichannel for DMA devices and the iLBX to handle high-speed access to main memory, the original Multibus interface (P1) is wide open for systems-level communications; for example, interprocessor communication that allows bus masters and slaves to communicate with fully defined interrupt signals and arbitration. Neither of the two extension buses has the arbitration logic necessary to support multiple masters; they are designed to increase the performance of a single "master" Multibus node.

System performance can be further enhanced by combining the Multichannel with the iLBX bus using the iSBC 580 board as the link. Acting as a talker and listener on the Multichannel, the iSBC 580 transfers data between the Multichannel and system memory via the iLBX bus, bypassing the Multibus. Data can then be transferred from the iLBX and memory to the Multibus for transfer to other parts of the system, allowing other system tasks to use the Multibus resources while high-speed I/O block transfers are taking place.

This set up can substantially improve system performance because it helps eliminate the speed discrepancies between the Multibus and the Multichannel. The iSBC 580's independence from the Multibus makes it suitable for applications that transfer large amounts of data in and out of a Multibus system, such as links between the Multibus and a host computer and mass storage, graphics display and high-speed data acquisition subsystems.

Bus hierarchy: With the iLBX handling high-speed local memory transfers and the Multichannel I/O handling direct-memory access burst transfers, the Multibus is freed for systems communications.

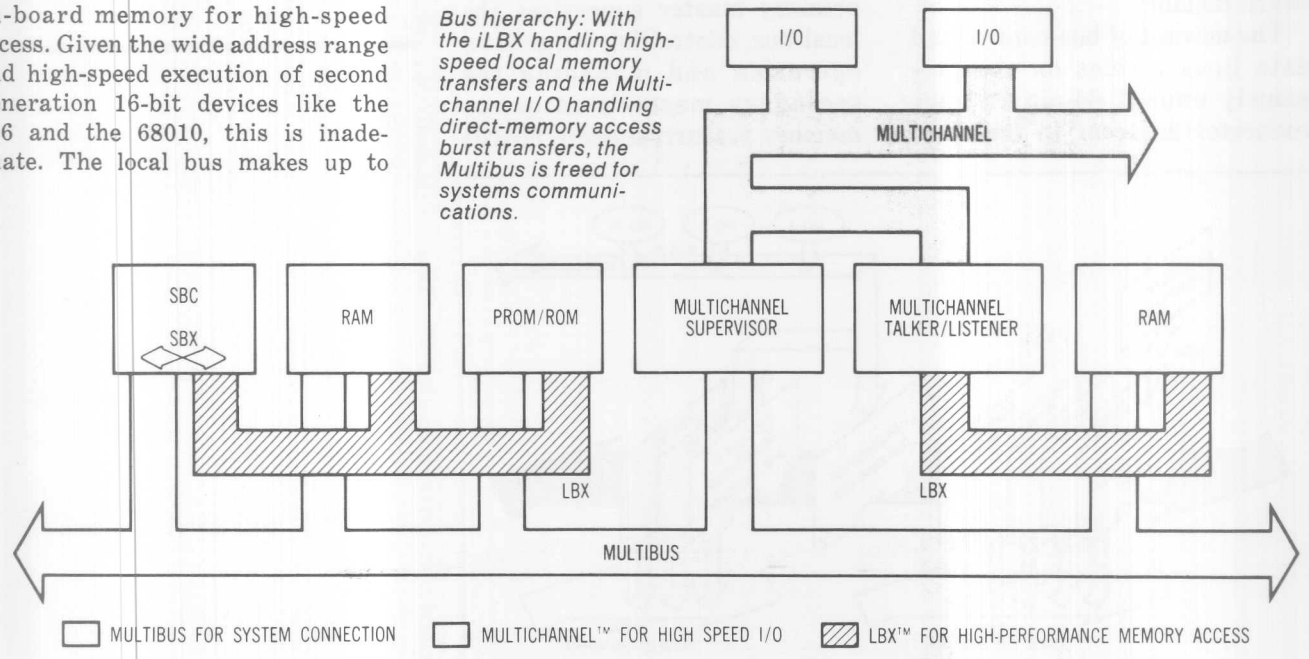

Computer Engineering

Intel Reveals Plans For 32-Bit Successor To 16-Bit Multibus

PORTLAND, OR — Intel Corp. has revealed plans for a 32-bit successor to its highly successful 16-bit Multibus. Multibus-2 is aimed at warding off any market erosion that the Motorola/Signetics/Philips 32-bit VME bus is making, as well as anticipating DEC's new 32-bit bus to be introduced this summer.

"We need a well-established 32-bit bus and the boards to support it when the iAPX-386 arrives," said marketing manager Rich Bader. It is believed, however, that the move is aimed more at recapturing market share lost to the VME than as an attempt to beat DEC to the punch.

Right now, the leading industrial board products are Multibus and DEC's Q-bus and Uni-bus. However, the increased functionality of the VME is presenting a challenge. And DEC is expected to introduce a new faster multiprocessor-oriented 32-bit bus along with its low-end micro-VAX, the 11/610, and high-end 11/810, sometime this summer.

Multibus-2 will be based on the same open-system concept on which the Multibus is based—all specifications will be made public and readily available. Although details are not yet being released, "It will be based on a hierarchical architecture akin to the LBX concept," Bader said. LBX refers to Intel's current use of the long unused second (P2) connector in the Multibus specification as a dedicated high-speed memory bus for the main processor. Careful attention is also being paid to multiprocessor configurations.

Putting those two concepts together suggests a multiconnector form factor that can be incrementally expanded to accommodate the special needs of new devices—perhaps even one that can make use of the existing 16-bit Multibus boards.

It is also rumored that Intel is working on a large-scale integrated circuit that will interface most devices to the bus. This is a component that is sorely lacking in the industry; often the circuitry to interface to a bus can consume as much area on a board as the actual board-function circuitry. A special interface chip that handles the timing and protocols would go a long way to making the Multibus-2 as successful in the 32-bit world as the Multibus is in the 16-bit world.

It is not known if DEC is working on any kind of LSI interface chip for its new VAX 11/610 or 11/810. But observers strongly believe that both will contain considerable new LSI, particularly the 610 that is termed by some to be based on a "VAX on a chip." It is more likely that the 610 will use some enhanced version of the J- or T-11 chips that emulate the PDP-11/70 and -11/23 instruction sets respectively.

The desktop VAX, code-named Scorpio, is said to include a two-board set using a new 32-bit bus.

CONTROL ENGINEERING

Local bus extension

The iLBX, local bus extension allows a microprocessor to logically address up to 16 Mbytes of local system memory. Designed as a high-speed, high-bandwidth processor to memory execution bus, the iLBX is for use with existing and planned 8- or 16-bit microprocessor-based single-board computers, as well as current and future high-density products.

INTEL CORP., Hillsboro, OR

ELECTRONIC DESIGN

Multibus extension

Intel (Santa Clara, Calif.) announced an extension to the Multibus architecture that allows a microprocessor to logically address up to 16 Mbytes of local system memory. The iLBX extension employs the P2 connector on the existing bus and permits transfers of 8-bit data at 9.5 Mbytes/s and 16-bit transfers at 19 Mbytes/s.

CANADIAN DATA SYSTEMS

SNA emulation software aimed at Multibus OEMs

Ottawa-based Xicom Technologies Corp. and Intel Corp. recently signed an agreement to bring Multibus-based OEM's and manufacturers full IBM communications capability.

According to Xicom, this will make it possible for micro manufacturers to develop products for use within the IBM Systems Network Architecture (SNA) market.

The SNA emulation software developed by Xicom uses Intel's iSBC 88/45 advanced datacomm processor Multibus board. The system, called SNA MicroNode, is said to allow users to have SNA networking capability without sacrificing their own data processing capabilities.

The MicroNode was recently introduced and Intel is providing product development, marketing and promotional assistance.

Multibus-based manufacturers will now have the opportunity to be among the first into the IBM terminal market requiring SNA compatibility, said Warner Sharkey, president, Xicom.

As 16-bit boards mature, 8-bit designs stand fast

Which board is right for the job?

Putting this information into perspective is James Johnston, applications engineer at Intel. He defines the two basic types of systems in slightly different terms. A multiprocessing system has "multiple processors that are homogeneous in that each processor is identical. Each μC board is dynamically allocated a task via a controller or a supervisor." On the other hand, he defines a multicomputing system as one that contains "heterogeneous μCs that might be identical or different; the system does not care. Each board must have the capability to stand alone with individual I/O, memory and software."

Johnston claims that "systems designers are now looking more to multicomputing techniques, ensuring continued viability for 8-bit boards. If throughput is not an issue and an application does not require large amounts of memory, an 8-bit board is always going to win out. The 16-bit implementation is going to cost twice as much in hardware as the 8-bit board—and be more difficult to design with as well."

The reason? The 8-bit boards have a huge body of well-tested assembly-level routines, operating systems, high-level languages and applications software that can usually suit your needs with little or no modifications. And when standard software is not available and custom development is required, the necessary programming tools and development systems are less expensive, easier to use and more widely available than those for 16-bit systems. Furthermore, more designers are experienced in 8-bit-system design.

EDN APRIL 14, 1983

intel®

ARTICLE
REPRINT

AR-268

April 1983

"Standard Buses Capture Fancy of Most OEM's" "Building Blocks for Micros"

Mitchell York
Michael Azzara
Computer Systems News

Intel Corporation, 1983
Order Number: 230727-001

Standard Buses Capture Fancy Of Most OEMs

By Mitchell York

Before OEMs or systems builders can choose a single-board computer around which to configure a multiboard system, they must first decide on a system bus.

It used to be they simply designed proprietary buses. But the evolution of board-level computers and standard bus architectures that began in the mid-1970s has changed all that: OEMs and systems houses now perceive the use of standard buses and bus-compatible board-level components as a way to reduce costs and lay the groundwork for the integration of future technological advances. The days of the proprietary bus are numbered, industry executives agree.

Several standard buses are now fighting for market share, with Intel Corp.'s Multibus leading the pack with more than half the 16-bit segment and one-third to one-half of the 8-bit market, according to managing analyst David Aronovitz of research house Gnostic Concepts Inc. Pro-Log Corp.'s STD bus is another major contender in the 8-bit arena, he said.

And both are bracing for stiff competition from a more recent offering sponsored by a coalition of semiconductor manufacturers led by Motorola Inc. The coalition's VME bus is compatible with 8-, 16-, and 32-bit architectures.

Additional competition is provided in the 8-bit market by the S-100 bus developed by MITS Inc. and in the 16-bit market by Digital Equipment Corp.'s Q-Bus. The proprietary bus is still an option, but industry experts are betting against it.

Many OEMs and systems builders are looking to standard buses to reduce engineering overhead and help position them to expand and modify their systems when new technologies become available. Those two factors — economics and technological advances — are the driving forces behind the mounting attention being focused on standard buses, according to industry executives.

"The technology has moved to a point where the user can integrate the baseline of his system onto a single-board computer," said Gary Sawyer, board products marketing manager for Intel's OEM Modules Group. "But we have very few customers out there who use a single-board solution. The primary advantage (of a single-board computer) is more content on a single board, so that as an OEM configures a system on his bus, he has a simpler configuration to perform. And that's why the bus question becomes so important," Sawyer said.

"The market, now that it has reached an adolescence, is maturing to a point where customers have to consider what bus they're deciding on first, and where that bus can take them in the future. Then they can evaluate the vendors and the products that are available.

While many OEMs still stand by proprietary buses, there are growing indications that the economics of the OEM business will not allow them to resist standardization for long. "There are still companies out there that are emphatically clinging to dedicated systems architectures for mini and microcomputers," said Ray Burkley, president of Astraea Computer Corp., Sunnyvale, Calif., which is marketing a VME bus-compatible single-board computer. "They are the ones that will end up getting hurt in the long run, because all the other reasons germane to using (standard) buses make sense."

"I think a proprietary bus is a definite tactical error," said Gnostic's Aronovitz. "It locks you out of the market, it forces you to design everything in-house, and it delays your entry into the marketplace, he said.

Standard Advantage

One reason for moving to standard architecture is that it is becoming prohibitive to design systems starting from the chip level, some OEMs said, unless the number of systems needed is high and unless the technology is likely to remain constant for a long time.

Observers say there is a proliferation of systems builders making small quantities of systems aimed at profitable vertical markets. Because these OEMs do not enjoy economies of scale and cannot afford to build from scratch, they are increasingly turning to standard architectures. In addition, with semiconductor and microprocessor technology moving so rapidly, it is likely that systems will have increasingly shorter marketing life cycles, thus further fueling the standard bus movement.

If systems builders are thinking about standard buses, vendors cannot be far behind. Aronovitz reports that there are 300 vendors in the board-level market and all have aligned their products with one or more of the standard buses.

The attention standard buses are receiving cannot help but grow exponentially in the coming years, industry executives say, because of backing from such powerhouse companies as Intel, DEC, and Motorola. Those companies and several others are helping create vast secondary markets to supply bus-compatible products geared for the systems builder. And the more products put out on the market for OEMs, the more OEMs are likely to use them and the buses they are made for, according to observers.

While OEMs and their suppliers may differ about which bus is best, there seems to be widespread agreement on the need to seriously address systems building starting at the board level and employing some standardized bus. It's a simple question of speed, according to Jeff Gorin, product manager of Motorola's MOS Integrated Circuits Group.

Faster Integration

"The fundamental advantage of modular product implementations using boards with common bus interfaces is that the systems integrator is going to be able to accomplish system integration faster — assuming he can find board-level components — by plugging together existing, documented boards rather than starting from scratch from chips," said Gorin.

Speed of product entry was the key factor that led Control Automation Inc., a Princeton, N.J.-based maker of factory automation equipment, to adopt an industry-standard bus. The company, which makes industrial robots and vision processors, was concerned with cutting down its engineering time and ramping up production, said director of electronic development Abe Abramovich.

Control Automation accordingly decided to build its robots with boards compatible with Multibus. "We came up to speed in our product development cycle more rapidly by buying off-the-shelf computer boards with all the functionality we require. By not having to design circuitry, we were able to apply our energy where it was really required — value-added interfaces."

The time savings that standard buses yield can be enormous, said Astraea's Burkley. Standard buses, with the myriad off-the-shelf board-level products available, allow an OEM to get to market from six to 18 months sooner than if the systems builder were to design his own proprietary bus and boards, Burkley said. "It's a make-buy decision. Do you want to take 30 months to develop a proprietary product, or buy at the board level, add software, and be in the marketplace?"

Once the decision is made to build with standard boards that are compatible with standard buses, the issue becomes which bus to use. Several have emerged as leaders because they

have won acceptance from components manufacturers who intend to create an array of products for systems builders.

One of the latest buses to achieve a growing degree of recognition is the VME bus. Motorola, Signetics Corp., Mostek Corp. and Thomson-EFCIS last year all threw their weight behind VME, and now an IEEE committee chaired by Signetics' engineering manager Michael Clader has been formed to draft a standard.

Proponents of VME, which is specially designed for use with the 68000 chip, say its main attribute is its ability to support 8-, 16-, and 32-bit microprocessors. This is important to OEMs that want to design 16-bit systems now that will work in the 32-bit environment when the technology is available.

"The VME bus is the only bus architecture that really allows us to build things today for both 16- and 32-bit," Burkley said. "It's our intention to build 32-bit supermini-computers. A lot of LSI components we need aren't there today, so we selected a bus architecture that will allow us to grow from the present 16-bit capability to 32-bit.

"We'll start building board-level products now that we can sell and which will generate feedback as well as profits. When full 32-bit capability is available, we'll only have to add a couple of boards to do the upgrade."

VME is also cited by systems integrators as more expandable than several other major bus options. "It gave me more real estate. That was critical to me," said Roger Vass, president of Victory Computer Systems Inc., San Jose, Calif. With VME, he said, Victory can achieve the cost efficiency of single-board computer design without sacrificing system expandability.

"VME has architectural aspects that are directly coincident with my marketing strategy," Vass continued, "I am going into the fault-tolerant, multiple-CPU, transaction processing market within a year." VME's speed is essential to succeeding in Victory's targeted markets, Vass said. "I need that kind of bus speed in order to execute fault-tolerant systems. I got it."

What Victory also got by adopting VME, Vass said, is a bus architecture that has been endorsed by major semiconductor houses, which gives it market credibility. "And it's a new bus. From a marketing standpoint, it's attractive because it's new."

VME is frequently compared to Intel's well-established Multibus. VME supporters maintain that Multibus can not fully support microprocessors as powerful as the 68000. Multibus enthusiasts counter by noting the constant stream of upgrades that have been made in the bus since its introduction to the OEM market in 1976.

Sawyer noted that even though Multibus was originally conceived as a bus that would address the 8-bit market, it has been adapted for 16-bit. "There have been six improvements made in Multibus to adapt to new VLSI. We are at the forefront of yet another evolution of Multibus to adapt it to our 80286 microprocessor."

LBX Refinement

The latest Multibus refinement is LBX — Local Bus Extension — which the company said allows a microprocessor to address up to 16 Mbytes of local system memory at very high speeds. LBX, Intel said, supports a data transfer rate of 9.5 Mbytes per second for 8-bit data and 19 Mbytes per second for 16-bit data. Intel said six single-board computer manufacturers have notified them of plans to develop products based on LBX.

It is just this kind of secondary market support that is crucial to the success of a bus, according to industry executives. There are now 165 suppliers of more than 1000 Multibus-compatible products, according to Fred Mazenac, president of Ironoak Co., the La Jolla, Calif.-based publisher of the *Multibus Buyers Guide*.

Motivating Factor

"Very clearly, one of the most important motivating factors in users selecting a bus is the number of vendors and products available for that particular bus," said Mazenac, who is now involved in assembling a Multibus manufacturers association.

Intel said the market last year for single-board computers was $375 million, about one-half of which was Multibus-compatible. By 1985, Intel said, revenues are expected to grow to $800 million, and Multibus is expected to retain its approximate 50 percent market share.

Sawyer said another boon to Multibus is its endorsement by the IEEE, which recently declared it a standard. With a firm, unchangeable standard in place, more vendors are expected to develop Multibus-compatible products, Sawyer said.

While Multibus is being used for a broad range of applications, Pro-Log's STD bus is geared primarily for industrial uses. Companies that build products compatible with the STD bus, also known as Standard bus, note that its small card size makes it attractive for single-board systems vendors that focus on industrial applications.

Dick Thomas, director for product development for Pro-Log, said STD "has an I/O intensity rather than a number-crunching intensity," which makes it better for control purposes than for data processing.

Thomas said the main advantage of STD is its price competitiveness. STD-based products are less expensive than Multibus-based products, but also less versatile. "What tends to happen (with the large Multibus card) is that people use it up and put a lot of things on the board just because there is a lot of room. It can be overkill.

"If you need what's on a Multibus card, like two I/O channels, it probably becomes more cost-effective. But that rarely happens, and that's where our advantage is. STD allows you to modularize to specific functions. Consequently there is a big price advantage," Thomas said.

While STD hasn't made its way into many office systems because of its data processing limitations, the older S-100 bus is most often used in small business systems. S-100, developed in 1975 for the Altair home computer, "is very flexible and modular," said Mark Garetz, president of CompuPro Systems, Oakland Airport, Calif.

CompuPro works with S-100, Garetz said, because it can be used with almost any processor. "It's not that hard to put various processors on it; it's not that processor specific."

Garetz said the cost/performance ratio of the S-100 "is unequaled by any of the other bus structures." Equivalent products based on Multibus can cost three to four times as much as S-100 products, he said. □

Building Blocks For Micros

As Systems Get Smaller, The Role Of Single-Board Computers Gets Bigger

By Michael Azzara

On the most basic level, a single-board computer is just what the name implies: the essential components of a computer system — a CPU, memory, and at least one I/O channel — on a single printed-circuit board. But the widespread assumption that it's a single board that constitutes a computer system is usually wrong.

So what's right? Single-board computer vendors agree that their products are malleable; they can be used alone in some applications, while in others they must be tied with additional boards to make a working system.

The fact is that use of single-board computers is pervasive in the microcomputer industry: Such boards are the building blocks of most modern microcomputer systems. The majority are manufactured from scratch by high-volume general-purpose hardware vendors, but more and more systems houses and OEMs in narrower vertical markets are buying single-board computers, usually from one of the approximately half-dozen semiconductor houses that analysts estimate control about 98 percent of the board-level products market.

The phrase single-board computer was coined by Intel Corp. in 1976 and has become virtually a generic term for microprocessor-based CPU boards. Many single-board computer vendors, however, point out that the term is somewhat misleading, since most of their OEMs actually include more than one board in their final products. But they agree nonetheless that there are important advantages to cramming the essential components of a system onto a single board rather than using three, four, or more boards strung together to make a computer system.

The primary benefits they cite are increases in performance and decreases in both cost and physical size, not necessarily in that order — and not all vendors agree on the relative importance of each.

On the flip side of the coin, single-board computers are a cost-effective solution for OEMs in only a narrow range of volumes. In very low-volume applications, an OEM may find a packaged system to be a better buy, while high-volume applications may call for an OEM to turn to component-level integration. Most of the applications that fall in that range are technical in nature, such as controlling scientific instruments or industrial machines.

Within that volume range — usually anywhere up to 1000 units during the life of a system — the advantages far outweigh the disadvantages, vendors say.

Performance is improved over multiboard systems because execution speed and reliability both increase when more work is performed within the confines of one board, rather than through communications on a system bus. Cost goes down because using one board to perform functions that once required three or more saves on board expenses. And because the final system can be made smaller than multiboard systems, doors are opened to applications that previously required smaller systems than could be built.

Besides those technical advantages, many vendors point out that buying single-board computers allows OEMs to concentrate their efforts on their own expertise in a particular applications area, rather than expending engineering talent to duplicate electronics that are already available in standard packages.

There are also the obvious benefits derived by the OEMs whose applications can be handled by one stand-alone single-board computer. These applications tend to be in the area of industrial automation and process control, such as the control of medical and scientific instruments.

All single-board computers, however, are not sold to OEMs. In fact, the market in 1982 was evenly split between the captive and OEM segments, according to managing analyst David Aronovitz of Gnostic Concepts Inc. He said, however, that he expects the captive percentage to grow as board manufacturers like Intel and Mostek Corp. begin selling more of their own systems. Adding to the captive segment are companies like NCR Corp., which makes boards for itself but doesn't sell in the OEM market.

Aronovitz said Intel shipped approximately 30,000 16-bit single-board computers in 1982 — about one-third to one-half of the 16-bit market. Specific figures for 8-bit single-board computers weren't available, but Aronovitz said volumes are much higher — in the hundreds of thousands.

In the OEM segment, most single-board computers are sold through the major electronics distributors such as Avnet Inc. and Schweber Electronics Inc. Smaller vendors sell through manufacturers representatives, Aronovitz said.

But despite their increasing popularity, single-board computers aren't a panacea. Before buying such a board, OEMs must consider several factors, primarily the questions of application and volume. OEMs must first consider whether the application is suited to a single-board solution. If it is, the OEM must then determine whether the solution promises to sell in volumes that fall into the range that makes board-level integration cost-effective.

How long it will take to bring a product to market must also be taken into account. If an OEM must react quickly to a closing market window, single-board computers provide a faster answer than in-house design, and buying packaged systems is a solution that's faster still.

"There are economies of scale at the various levels of integration," said Gary Sawyer, board products marketing manager for Intel's OEM Modules Operation. "Components are the lowest cost, highest risk, and take the longest time to (bring a product to) market. At the other end are systems where the OEM's value-added is software. They're the highest cost, lowest risk, and take the least time to market.

"At the board level, we're right in the middle," Sawyer said. "Primarily, we're talking about a board crammed with technology, about 81 inches square, that goes into a laboratory or factory environment," he said. About 80 percent of Intel's single-board computers are used in those types of technical applications, while applications for the remainder are varied.

Sawyer said microprocessors are pervasive in the desktop computer area, but because the volumes for those applications are enormous, manufacturers primarily use components and design their own boards.

More Than One

Sawyer admitted that despite the name single-board computer, most OEMs use more than one board.

"But there is a great benefit to solving your problems within the confines of a single board and using your bus as a system resource. You move questions and answers across the system, on the bus, but keep real hard work off the bus. When you have hundreds of thousands of bytes of memory in a single-board computer and can do the work on such a localized basis, your performance benefits are enormous," Sawyer said.

"Another thing is that with several boards, you have to continually go back and forth through the bus to fetch instructions, move

memory, etc. What you want to do, in my view, is localize those functions and maintain that precious system bus activity for system-kind of functions, like communications between the CPU and disk drives or CRTs. You don't want to use that bus to move bits and bytes around like that," concluded Sawyer.

Even as single-board computer vendors tout their latest semiconductor advances, waiting offstage is the next major step in the miniaturization process — single-chip computers.

Semiconductor houses are already selling what they call microcontrollers — chips that include not only a microprocessor, but nonvolatile memory and I/O capability as well — in large volumes for simple control applications. Industry observers believe it's just a matter of time before chip makers are able to squeeze enough sophistication into a silicon wafer to make it indistinguishable from current single-board computers.

"That's the new wave," said Rod Zwonitzer, system products marketing manager for Mostek Corp. "Taking what you have on a single board and shrinking it onto a chip. This technology is a continually shrinking thing."

Existing microcontrollers, with most software embedded in ROM, are used to control predetermined functions in peripheral devices such as disk drives, CRTs, printers and keyboards, according to Graham Alcott, Intel Corp.'s microcontroller products marketing manager.

They are also used in much higher volumes in consumer products such as televisions, telephones, automobiles and microwave ovens.

Extremely Small

Because of their extremely small size and price, microcontrollers can cost-effectively replace mechanical and electromechanical devices in those applications, said Joseph Baranowski, assistant marketing manager of Intel's microcontroller operation.

And because those are such high-volume applications, microcontroller production is high. According to a recent Dataquest Inc. market report, U.S. semiconductor houses shipped approximately 50 million microcontroller chips in the third quarter of 1982, Baranowski said. He added that Intel will not even entertain orders for fewer than 1000 pieces.

Intel's Latest

Intel's latest device, the 8096, packs in a 16-bit CPU, 8k bytes of ROM, 256 bytes of RAM, and multiple I/O channels, including eight analog-to-digital converters, he said.

The analog communications capability is an important aspect of microcontrollers, since it enables them to control the functions of devices in a realtime environment, Alcott said.

Compared with Intel's first single-board computer, the 8010, which was based on the 8-bit 8080 microprocessor chip, the 8096 single-chip computer provides double the memory and more than double the performance, according to Baranowski.

"A lot of what we're now doing on single chips were actually board-level products less than five years ago," he said.

Future Trend

"The trend in the future is that microcontrollers will continue to be a dominant way people design," said Baranowski.

"The typical office of the future will have numerous disks, keyboards, printers, all with one or more microcontrollers, and there will likely be only one microprocessor acting as CPU for the whole system," he said.

Besides Intel and Mostek, prominent microcontroller manufacturers in the United States include Motorola Inc. and Zilog Inc.

Adaptable Micros Expand Capabilities To Provide Better Control Functions

Microprocessors have influenced the way we perform control functions. They, in turn, have been influenced by their performance in many different control functions. During the past year, micro manufacturers have developed more sophisticated and specialized machines to perform increasingly complex operations. Some of these "wunderkind" have been optimized to more efficiently handle basic control problems, while others have been developed to answer needs created by microcomputers themselves.

ALAN J. LADUZINSKY, CONTROL ENGINEERING

The evolution of microprocessors and other digital chips has continued at a rapid pace. Microprocessors continue to progress and branch off along the road to new and different control functions. Manufacturers have accelerated this evolution by providing a flexibility that allows design engineers to develop chips that are more suitable than the old standbys for specific jobs. These semicustom chips combine digital logic and simplicity of function to emerge into measurement and control functions that were not previously possible.

Perhaps the most interesting development since last April's article is the lateral leap microprocessors have taken to the lower power demands of CMOS technology. Harris Corp.'s Semiconductor Group (Melbourne, FL) will introduce a CMOS version of Intel's 8086 with support peripheral chips by June of this year.

A precursor of CMOS MPU's is Motorola's MC146875G2, which is touted as the world's first microcomputer to combine CMOS and EPROM technology. Additionally last September, RCA agreed to develop CMOS versions of some of Motorola's microprocessor chips. Texas Instruments has also developed CMOS version of their micros.

The increase in performance and operational complexity that control engineers have put upon micros has forced a growth in chip functionality and the ability to access more memory. Microcomputers and computer boards have literally had to grow or develop new data paths or buses and structures to meet the control engineers expected performance levels.

The results of these different developments combine to provide digital devices that are more suitable in the control environment. They also allow the crystalline critters to carry or deliver digital functions to new applications.

The adaptability of microprocessors can be chameleon-like in nature. Much as the chameleon can change colors to blend in with the background, so can the micro by changing to CMOS without changing function. The use of CMOS technology reduces power requirements for design, and thus weight and size. It is also less sensitive to electrical noise found in the industrial environment. This most remarkable adaptation allows control engineers to expand microprocessor control to many applications requiring smaller devices.

Market forecasts indicate that CMOS VLSI devices will experience one of the highest growth rates among the VLSI technologies this decade. This fact has not been lost on silicon designers. Several technology exchange agreements have resulted.

In an agreement announced in June of 1981, Harris Semiconductor (Melbourne, FL) and Intel agreed to a technology exchange for the development of complementary metal oxide semiconductor (CMOS) microprocessors and peripherals using Intel's designs. Intel would be the second source of the Harris CMOS chips. In June of this year Harris will introduce a CMOS version of the 8086 (80C86), along with all the supporting peripheral chips.

The 80C86 and its peripheral sup-

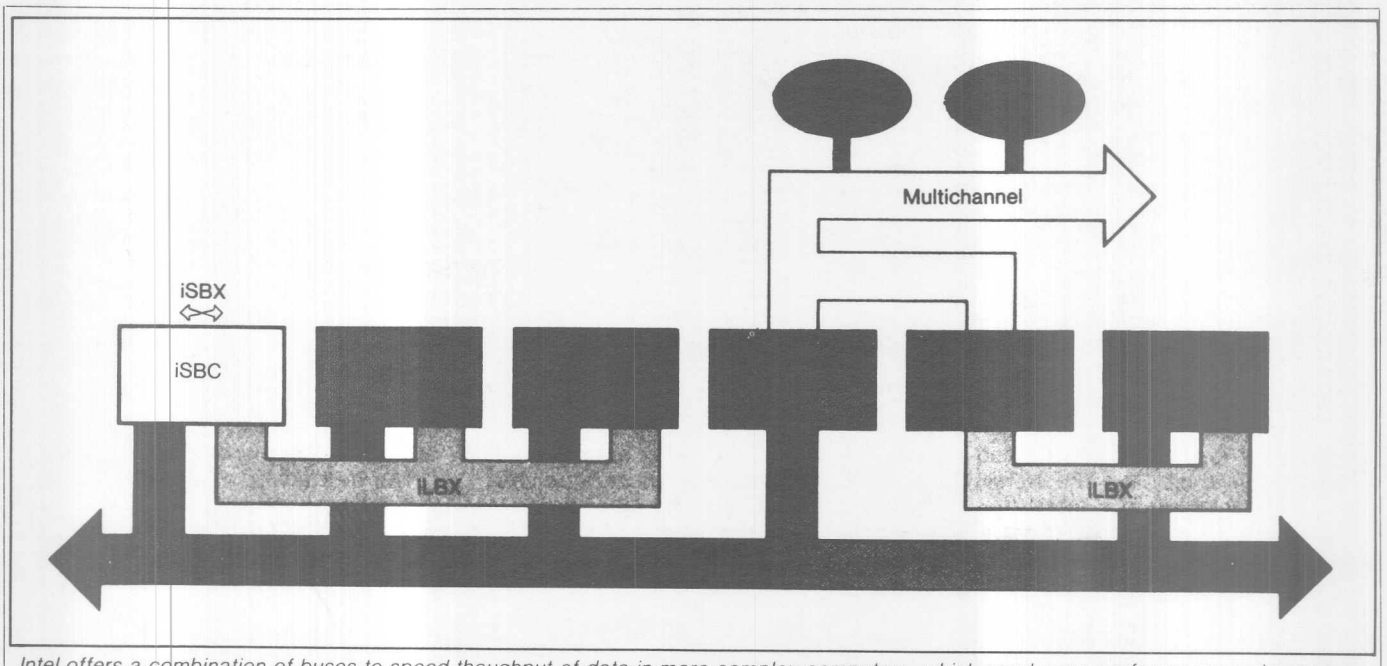

Intel offers a combination of buses to speed thoughput of data in more complex computers, which need more performance and memory access such as the 80286. A total of four different buses include the traditional Multibus, the iSBX for on-board expansion, Multichannel for high speed I/O, and the new iLBX for faster memory access.

230727-001

port are pin compatible replacements for their NMOS counterparts. It will be TTL compatible with capability for retrofit. The basic advantage to Harris is the existing software base for Intel micros. Intel gets a low power consumption micro that features, Harris states, no performance degradation. The 80C86 chip uses active terminators instead of pull up resistors to reduce power drain.

The 80C86 will come in both 5 MHz, and 8 MHz versions. An additional advantage will be the static capabilities of the chip system which will allow the system to stop and operate in a dc (static) mode. The static design also allows single step system debugging, while NMOS chips have to operate at 2 MHz. The 82C59A priority interrupt controller can be employed to activate the device with external interrupts.

Texas Instruments (Dallas, TX) has also developed a series of CMOS 8-bit single-chip MPU's this year. The CMOS TMS70C00 and TMS70C20 employ the same chacteristics of their NMOS counterparts. The -C20 has 128 bytes of RAM, and 2 kbytes of on-chip ROM.

Do it yourself micros

American Microsystems, Inc. (Santa Clara, CA) takes adaptability of microcomputers literally. Using the term AMU for alterable microcomputer, the company combines computer aided design with a library of standard and custom functional cells that allows a system designer to build a special purpose VLSI chip. This VLSI answer to a design problem might otherwise require many standard chips or a custom VLSI circuit. In both cases development costs could run many times more than the semicustom AMU chip. Development cost for the AMU runs from $30,000 to $80,000.

Conceptually the AMU system allows an individual who is not familiar with digital circuit design to specify in a system of functional block form the type of VLSI circuit needed to develop a product. Each of the blocks can be specified individually and called out separately. This way, the designer can fashion custom functional cells one at a time to simplify design and testing. Functional cell integration is also simplified by a standard cell topology.

The work is done in CMOS using a 9900 family 16-bit CPU. The standard cell library includes I/O, counter/timer, clock, RAM, ROM, output pad driver, address and bus driver, and serial communications. A customer can develop custom cells on a proprietary basis, or AMI will split the development cost of a custom design and add it to its standard cell library.

The semicustom design approach is said to have a typical break-even point of 20,000 units. Under special circuit design circumstances the break-even point can be as low as 5,000. Applications include a portable oscilloscope, and A/D and D/A converters with a microprocessor on-chip. (Linear circuit manufacturers have been working toward this from another direction. See CE, Dec. '82, pp 40-42.)

Savings also can be influenced by factors other than the cost of chips and design. The use of CMOS design permits a saving on the power required, and thus the size of the supply. This in turn permits saving in housing size, and reduces or eliminates the need for cooling and heating because of a wider operating temperature range. Applications cited so far point to portable instruments which require small size, low weight and power consumption, and high functionality. The reduced part count of a semicustom chip can also translate into increased reliability.

Adaptability of another kind

Form determining function and adaptability in control was approached from a different direction by Texas Instruments and Seeq Technology (San Jose, CA). They agreed to cross license technology to put nonvolatile E²ROM on the TMS7000 family of microcomputer chips. The addition of E²ROM on-chip allows the micro to be programmed or reprogrammed on site electronically. This in effect can make the already very adaptable microcomputer even more so when reprogramming becomes necessary with changing job assignments.

The first device in this series, the 72720 *adaptive* microcomputer, will have the instruction set of the original TI unit, with a new command which will allow the 72720 to alter and progam its nonvolatile memory. This permits a change of programming with the press of a button. Such capability will allow a programmable machining center, robot, or other part of the automated factory to easily change programming as new tasks arise.

The 72720 will also include: expanded RAM to a total of 256 bytes to facilitate scaling; additional registers and control logic to provide a security feature that inhibits external access to the internal memory once the program has been verified; and Silicon Signature (to establish programming algorithms and voltages) and DiTrace (128 E²ROM that carries detailed processing and test information) features.

Form makes function

The complexity of microprocessors provides more functionality to the designer. Not only does the micro evolve in control applications but so does the engineer and designer who applies it. A familiarity with microprocessor control breeds learning and understanding, and nefarious thoughts of more sophisticated applications and higher performance. While which came first, the need for more performance or more complex microprocessor circuitry, may not be answerable, both are an inevitable trend.

As more circuits are placed on a single chip, a change occurs in the way it is programmed. Software functions are transformed into hardware thus simplifying programming, and allowing the MPU to perform tasks quicker. System reliability can also benefit when ROM and RAM join the CPU on the same chip. The different combinations of function provide control engineers with a plethora of solutions for control problems.

The traditional way of building standard microprocessors from VLSI chips in silicon in large numbers still accounts for the bulk of micro shipments today. This is explained in part by the fact that many requirements of control can be categorized into a limited number of groups. So a standard micro can still fill the bill for many control needs.

Intel Inc. has introduced a series of VLSI chips that combine many functions. The chips include the 8096 series controller and the 80186 and 80286 microcomputer families. This group was designed with a systems approach in mind. The 8096 and 8396 were designed specifically for use in control.

The company believes the trend to more sophisticated and more complex operations in control necessitates placing more functions on a chip. One can also add to this the increasing demand for performance and memory. This necessitates different functional organization on a chip. The 8096 series is a new generation of microcontroller which incorporates a new organization with a "flat" architecture for direct memory addressing as compared to the 8051's segmented memory address. This results in faster memory access and the ability to address large amounts of outside memory. The 8096 also adds new I/O and communications capabilities on a chip to address the need for higher functionality.

The 8096 incorporates an I/O system that relieves the central processor of some of the overhead and so speeds operation. To make the chip look more like a control system, an eight-channel 10-bit A/D converter shares space on the chip along with a watchdog timer, an eight-level priority interrupt structure, PWM output, full duplex serial I/O, and 40 parallel I/O ports.

230727-001

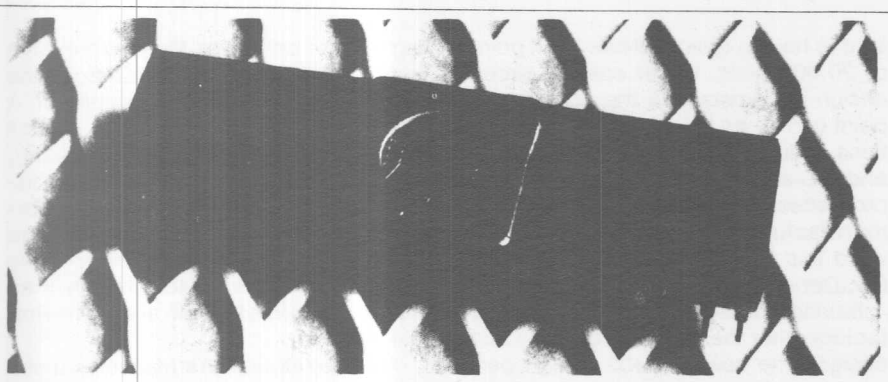

Motorola's 68000 series microprocessors deliver different combinations of ROM/RAM to support user needs. Direct memory addressing capabilities speed memory access.

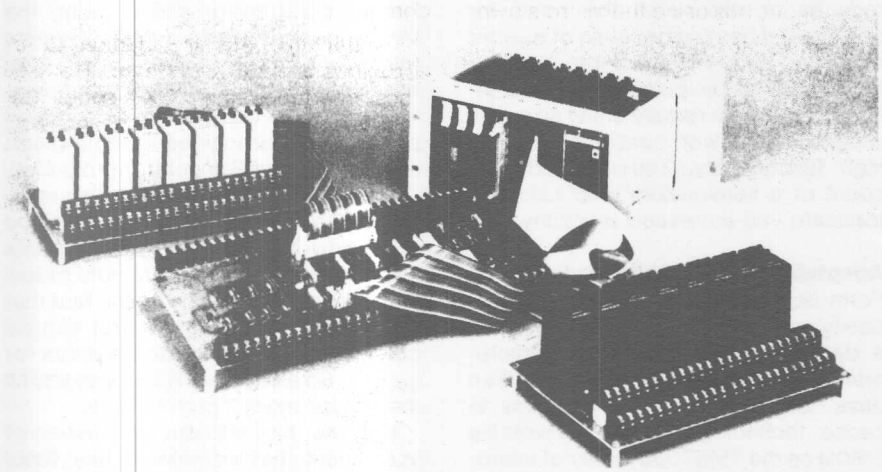

National Semiconductor's CIM system has DIN 4162 socket connectors and a synchronous CIMBUS structure for control applications. Here the system interfaces to solid state relays.

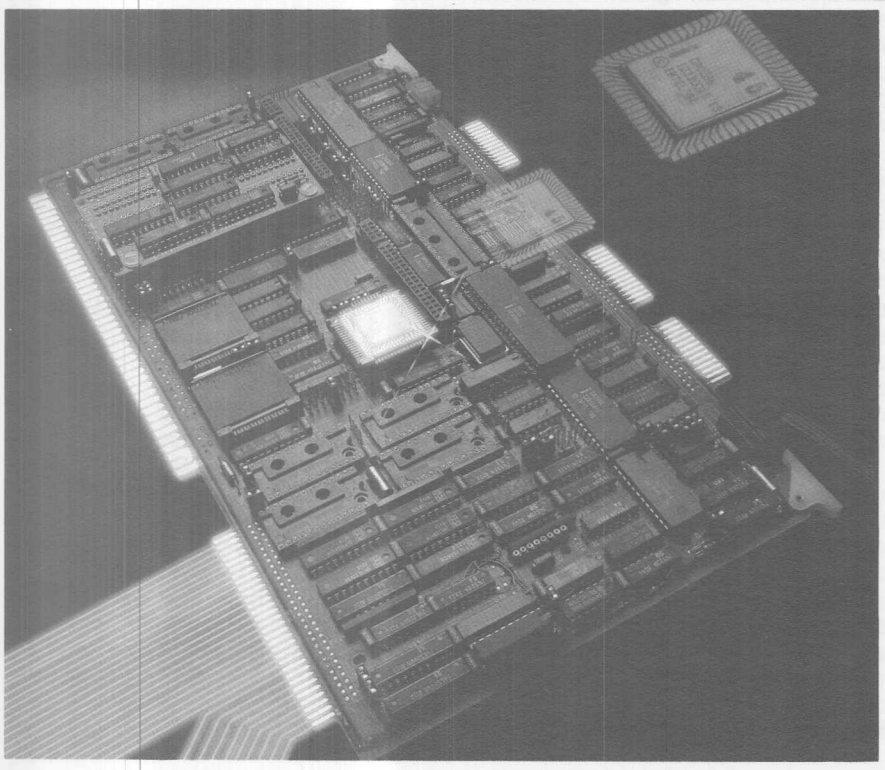

Intel combines an advanced microprocessor structure and multiple bus structure to answer the control engineer's increasing needs.

Addressing adaptability, the 8096 and its 8051 forebearer can be modified with special custom code to perform a particular operation on the plant floor. During an interview with company employees, it was learned that a modification of the 8096 optimized for robotics algorithms may make its debut in the future. As an aside, future generations of the microcontroller family will be developed in CMOS.

The 80186 and 80286 are two chips designed to work with the 8096 as well as by themselves. All of these VLSI chips have hardware functions that may be transparent to the user. Transparency enables the engineer to design a control system without needing to know all the intricacies of the chip. This user friendliness is brought about because of the explosive nature of industry's tendency to use a good thing faster than engineers can learn how to support it.

Intel promotes this "-86" group of chips for the automated factory by pointing to their communication abilities, thus forming a system at the chip level. The 80286 is endowed with on-chip memory protection which in turn permits 4 μsec response time. The micro has a "pipelined" architecture that gives it the ability to identify invalid op codes before execution, and also check out protection qualifiers without degrading performance.

At the 8-bit level Motorola has adapted its 6805 microprocessor family's memory addressing. The resulting micros can be modified to work with ROM/EPROM sizes in a range from 1 k to 3.8 kbytes, along with 64 or 112 bytes of RAM. The ROM/EPROM to RAM ratio can be any combination up to the 64 kbyte addressing capability. In addition to special I/O, and 8-bit A/D converters, unique features include a "self-check" test mode in on-chip ROM to check all features including I/O and RAM, and thus reduce test time on incoming quality control. The 6805 family also features a bootstrap ROM allowing the micro to program its own EPROM to eliminate programming equipment.

A matter of form

Charles Farmer, Director of Engineering for Honeywell Process Management Systems Division (Phoenix) says it helps to have the right form or structure from the start. The company's CP1600 microprocessor has changed only by increasing the amount of addressable memory since its introduction. This was done in software by using a paging technique to increase memory addressing to 256 k. The software technique is slower than the hardware incorporated in other systems but

the company still considers it to be adequate.

Farmer points out that if the user gets the opportunity to address more memory, he will. Honeywell's philosophy leads to market driven adaptations of microprocessors for use in control. So far this has not led to an increase in functionality at the chip level that is apparent in the new developments of companies who might be more oriented to OEMs. The main difference for this thinking probably relates to the fact that the Honeywell chips are used in the company's own control products.

Functional modifications

Rockwell International (New Port Beach, CA) adds several features to its R6500 microcomputers to satisfy customers. In addition to ROM options, the company provides four timer/counter modes of operation, a timeout (or watchdog timer) on-chip, and five different types of interrupts. They include edge-sensitive lines (both positive and negative), non-maskable, counter underflow, serial data received, and serial data transmitted.

Additional functions available on the R6500 family include divide by two and divide by four counters. The microprocessor can also be configured to act as master or as slave MPU's.

Of boards and buses

The growth of function in microchips continues, but a single chip does not a computer system make. While many microprocessors have been developed and used for control, they must all work within bus structures.

As users require higher performance and more sophistication, more memory addressing comes into play. Intel forecasts a faster growth rate for 16 bit system purchases as compared to estimates for 8-bit boards. A 16-bit system addresses both the needs for higher performance and greater memory addressing. However, while faster execution of commands and memory access is possible on chip, the limiting factor becomes the data buses which microprocessors must use to interact with other components of the microcomputer system and the outside world.

National Semiconductor's (Santa Clara, CA) Steve McGinness, Product Marketing Manager of board-level products, cites three factors that determine real time performance. They are the number of interrupts, the performance of the computing element, and availability of a multi-tasking software environment. While the third factor would require an article in itself, the number of interrupts and computing performance can be readily discussed.

McGinness points out that his CIM boards have 12 vectored interrupts (eight on the CIMBUS, and four additional on the CPU board). This compares to eight on Intel's Multibus and the S100 large form factor boards. The STD bus offers only one. The number of interrupts needed in a system would help determine the type of bus needed.

The performance of the CPU is affected by more than just clock speed. National Semiconductor's CIMBUS is a synchronous bus, like the STD bus, with many of the bus signals being buffered versions of the microprocessor's own bus. McGinness points out that a synchronous bus using the micro's own clock will never outrun the rest of the system regardless of speed. He also states that an asynchronous bus will always have a usable bandwidth somewhat less than its bus clock, due to sychronization and arbitration delays.—"The Multibus, for example, has a 10 MHz bus clock, but a usable bandwidth of around 3 MHz."

National uses the DIN 41612 pin-in-socket connectors in board level products. McGinness states that as a result of this connection system, the CIMBUS can be enhanced from 8-bit architecture to 16 and even 32-bit architecture without changing the form factor. On the other hand, board/bus architectures using the card-edge connectors cannot increase the number of bus signals without changing the form factor meaning; the connector, the card cage, and so on.

More buses

Motorola (Phoenix, AZ) uses a combination of buses to speed signals. The company's VMEbus can be considered as a subset of the larger form VERSAbus. The difference being the number of priority arbitration levels. The VMEbus has four levels of priority arbitration versus five levels for the VERSAbus. The main difference between the two relates to form factor and the amount of memory needed to perform—VERSAbus can support more memory. Both offer asychronous operation at a bandwidth of 20 Mb/sec. Literature indicates that a mix of 8, 16, and 32-bit processors can be mixed using a common backplane.

In addition to a choice of bus structures, both VERSA- and VME- modules offer an additional bus called the I/O channel. This new system architecture permits modular I/O expansion on the local processor bus. This frees the VMEbus, for instance, to handle simultaneous high-speed data exchange and multiprocessor access requirements, while permitting most lower speed I/O activity to take place through the local I/O Channel. The I/O Channel is 12 bits wide. It allows 8-bit parallel bidirectional asychronous communication 2 Mb/sec throughput.

Still more buses

A different approach to moving information around the board involves the use of several buses. Intel Corp. has developed a system of multiple data paths to take advantage of its chip development efforts. These buses also establish a distinction among the different types of digital signals in a board level system.

Incorporated on the iSBC 286/10 computer board are the iLBX (Local Board Extender) bus, and the "Multichannel" I/O bus in addition to the Multibus and iSBX systems. The iLBX bus interfaces the "CX" series RAM cards directly through the P2 connector to the iSBC card. This allows the CPU board to address memory expander boards as though they were on the same card. The iLBX bus provides 16 Mbyte addressing, and a bandwidth of 9.5 Mb/sec for 8-bit words, and 19 Mb/sec for 16-bit words.

The "Multichannel" I/O bus provides Intel's board level products with a fast data path between the microcomputer and as many as 15 block transfer devices. The bus consists of 16 address/data lines, six control lines, two interrupt lines, parity lines and reset. Maximum throughput is 5.3 Mb/sec.

These multiple path data buses can also interact with each other. Multichannel is interfaced to the Multibus system bus through the iSBC 589 intelligent I/O controller. The iSBC 580 Multichannel to iLBX bus interface allows I/O devices access to RAM.

Providing just enough

RCA is a company that has been making microcomputers in CMOS for some time. The company's stable has many horses, including the company's 1802, CMOS versions of some of Motorola's micros, and the RCA version of a Z80 that should be out now. The computer is configured on modular boards in which CPU, memory and I/O can be on a single board. Over 80 modules can be combined to fit user needs.

Company staff pointed to the fact that such systems can deliver just enough computing power to perform the desired task. This saves money on installation when a more completely bundled system is an alternative. The RCA MICL software language allows the control engineer user the ability to program the company's system in a high level language much like a ladder diagram. Other High Level Languages also available are BASIC, and FORTH. □

230727-001

Data Files

By Phil Koopman
& Colin Johnson

MULTIBUS EXPANDS TO 16 AND 32 BITS

EE Times has learned that **Intel** is about to pop a new 16-bit version of the Multibus called the iLBX, for Local Bus Extension. The first public look will be at this year's Comdex Show in Las Vegas. Since its inception in 1978, the Multibus has undergone a steady growth in capability. And Intel's *de facto* control of bus specifications has pretty much ensured that older cards will work with bus upgrades. As an example, Intel's original 8080-based microprocessor card will still work with current versions of the Multibus. But there's still the problem of 16- and 32-bit data paths and the ever-increasing speed of microprocessors.

Expect Intel to show off its 16-bit version of the Multibus with a datatransfer rate of 20 MHz. The 16-bit bus will use the P2 connector. The architecture of the 16-bit bus structure will be similar to that used with microprocessor cards that had a dual-port memory—in effect, the 16-bit bus will be a processor-to-memory setup. Provisions won't be made to handle communications with peripheral cards, like those for analog I/O—an ability that was sacrificed to gain the 20-MHz transfer rate. Cards other than CPUs and memory will still access common control and data via the regular Multibus data paths on connector P1.

EE Times has also learned from another source that Intel has a 32-bit version of the Multibus setup in the works—possibly to be called the Multibus II. But don't expect this version until sometime next year. In fact there may be two proposals put on the table for the 32-bit configuration. Most likely Intel's setup will win out. Intel has been putting a market push behind the "new" Multibus. It hosted a breakfast about three weeks ago at a local Multibus show in Detroit, where it handed out a marketing overview—sans detail tech specs—to Multibus vendors.

Like the 16-bit Multibus add-on, the 32-bit bus will use the P2 connector. To solve the lack of pins, Intel is planning to add more pins to the connector, jacking the pin count for the two-sided P2 connector from its present 60 pins to a total of 72 pins with the addition of three more pin widths to the connector's length. By adding the new pins to the upper end of the connector, downward compatibility will be maintained with the new 16-bit iLBX version.

—P.K.

Intel's Multibus Set For Official IEEE Approval

NEW YORK — After nearly two years of deliberations, IEEE will formally bless Intel's Multibus as IEEE Standard 796 by the middle of December.

Rich Boberg, chairman of the IEEE's 796 Working Group, said the basic Multibus Standard, which defines the P1 connector, has been passed by the working group and by the IEEE Microcomputer Standards Subcommittee.

IEEE Finally Ready To Give Its Formal Stamp Of Approval To Intel's Multibus Next Month

(Continued from Page C1)

The proposed 796 Standard has also been voted on by the Computer Standards Committee, but there were two abstentions. The proposal has been submitted to the IEEE Standards Board in New York.

It is expected, providing the two abstaining voters can be satisfied, that the Standards Board will give its stamp of approval during its meeting next month. All that remains after formal approval is to print up that final 796 Standard, which should be done by January.

The present 796 Standard before the IEEE Standards Board covers only the basic Multibus and defines the P1 connector. The P2 expansion connector—slated for use as the iLBX 16-bit and the Multibus II 32-bit data bus (see "Data Files," this issue on page C2)—isn't covered in the present proposal.

"First P1, then the enhancements," said Boberg. The IEEE's approval of the Intel-originated Multibus is an important factor for the 122 vendors of Multibus gear. It signals that Multibus has reached a high level of maturity, both with users and in the marketplace. At present, there are more than 800 Multibus products on the market. Formal standardization ensures that a system designer will have compatibility as the architecture of the Multibus evolves in the future.

Cooperative Makers

Led by Intel's proposal to numerous Multibus makers, a Multibus manufacturers' group is in the middle of being formed. Open to all makers, the group will have its organizational meeting the middle of this month. Formal organization is slated for the Comdex convention in Las Vegas.

Preliminary plans for the new maker's group call for the formation of a working group that will thrash out how the 32-bit Multibus II system architecture will evolve. This group will also be headed, at least for starters, by Boberg. Boberg said that if any users or makers want more information about the new Multibus manufacturers' group or the working group for Multibus II, he will be pleased to help.

Boberg can be reached at Microbar Systems, Inc., 1121 San Antonio Rd, Palo Alto, CA 94303.

Microcomputers Boost Typesetting

Intel Corporation's single-board computer provides the brains for the Autologic APS Micro-5

FOR THE PAST 20 YEARS, typographers have been involved in an industry-wide changeover from hot metal to phototypesetting. This transformation has involved several distinct stages, or "generations" of equipment. For example, no sooner had the old-fashioned linecasters given way to photomechanical typesetters than the more efficient digital cathode ray tube typesetters made their appearance on the printing scene.

For several years, the new minicomputer-based CRT typesetters were too expensive for any but the large-volume publishers and typographers. By 1979, however, progress in microcircuitry allowed Autologic, Inc., to introduce its APS-Micro 5 CRT digital unit to address the financial and typesetting needs of the small- to medium-sized book, magazine, and newspaper publishers.

High speed

The APS-Micro 5 sets type at the rate of more than 1,000 lines per minute, producing very sharp type (resolution: more than 3,600 lines per inch). It fills the gap between Autologic's more expensive, large-volume digital typesetter that sets type at the rate of 4,000 lpm and the slower "second-generation" photomechanical typesetter, which sets type at the rate of only 100 lpm.

The APS-Micro 5 operates with little operator intervention (unlike its photomechanical predecessor, which required the constant attention of an operator). The Micro 5 operates easily—a few buttons run the entire typesetting job. Fitted with special computer programs, it can set automatically in formats varying from paginated books to newspapers to complex manuals. Thus, by replacing the complex mechanisms of the photomechanical typesetter with state-of-the-art electronic circuitry, Autologic has increased productivity of valuable staff and has streamlined the typesetting process.

Whether used for newspaper, magazine, or book production, the APS-Micro 5 accepts digitized input from most front-end computer systems in almost any form: magnetic tape, paper tape, or on-line.

Edited and processed copy is sent from the main computer or on-line entry device to the typesetter in digitized form specifying type style, point size, and space between the characters and between the lines.

The digital circuitry in the Micro 5 prepares the text, manages the data, and coordinates the various commands at an input rate of 10,000 characters per second. It selects the specified character from a library of fonts stored on a fixed disk in the system and quickly readies them for imaging.

Analog circuitry is then responsible for converting the digitized information into lines of print on a cathode-ray tube. Entire lines of print are beamed from the CRT through a fixed lens onto photographically sensitive, resin-coated paper or film.

Sophisticated electronic controls synchronize the movement of the photographically sensitive paper with the movement of the type image on the CRT tube. By moving the film or paper during imaging, the Micro 5 cuts down on typesetting time.

The typeset output is then pasted up on boards according to the specifications of the composer or the graphic artist. The camera-ready page goes to the camera room for the production of the final plate for the presses.

Not only has the typesetting process has been shortened by replacing the numerous moving parts of the old photomechanical typesetter with electronic circuitry, the new equipment is far more reliable because it has only two moving parts: the disk drive and the film or paper transport

Small machine

The supplier was able to design large-scale capabilities into a small (3x5½-ft) machine by using Intel Corporation's 86/12A single-board computer. The 86/12A condenses on a single board what was formerly a huge amount of electronic circuitry contained in a minicomputer.

The 86/12A is a prime example of how advances in electronic miniaturization have brought technical improvements to the printing field. It gives system designers like Autologic a fully developed computer building block at a much lower price than a minicomputer. It consists of the powerful Intel 16-bit 8086 microprocessor (see GAM, May 1982), random-access memory, and read-only memory on a single board.

"In developing the APS-Micro 5, our criteria were to design a fast, highly reliable typesetter at a lower cost than our large-volume, minicomputer-based system. Intel's single-board computer answered all of our requirements," says Tom LeJeune, an Autologic product manager.

"In choosing a computer on a board we looked not only at the board but also at the support system that comes with it," adds Elliot Freeman, vice-president–engineering.

Programming language

Program development for the 86/12A is aided by many support products such as the Intellec Series II Microcomputer Development System and PL/M-86 software package, Intel's high-level programming language.

"Intel's PL/M programming language simplified writing the software for the Micro 5. Intel's support systems dramatically reduced development time of the system," says Freeman.

Autologic has been developing large digital phototypesetters since the mid-1960s. Using the engineering knowledge gained from its minicomputer-based digital typesetters, Autologic started designing the APS-Micro 5 in 1977. Since its introduction to the market two years later, more than 350 APS-Micro 5s have been purchased by medium-sized book, magazine, and newspaper publishers, as well as by in-house typesetting operations. #

230727-001

Acknowledgements

We greatfully acknowledge the kindness of authors and publishers in giving permission to reproduce their materials in *How To Select a Multibus Architecture*.

1. **Canadian Data Systems, pg. 33**

   Reprinted with permission from *Canadian Data Systems*, April 1983 issue.

2. **Computer Systems News, pg. 35**

   Copyright © 1983 by CMP Publications, Inc. 111 East Shore Road, Manhasset, N.Y. 11030. Reprinted with permission from *Computer Systems News*.

3. **Control Engineering, pg. 33, pg. 40**

   Reprinted with permission from *Control Engineering*, Barrington, IL., April 1983.

4. **Datamation, pg. 5**

   Reprinted with permission from *Datamation*, April 1983 issue.

5. **Digital Design, pg. 7**

   Reprinted with permission from *Digital Design*, February 1983 issue.

6. **EDN, p. 34**

   Copyright © 1983 by Cahners Publishing Company. Excerpt reprinted with permission from *EDN*, April 14, 1983 issue.

7. **Electronic Business, p. 4**

   Copyright © 1983 by Cahners Publishing Company. Reprinted with permission from *Electronic Business*, April 1983.

8. **Electronic Design, p. 31, p. 33**

   Reprinted with permission from *Electronic Design*; February 17, 1983; March 31, 1983 issues.

9. **Electronic Engineering Times, p. 33, p. 44, p. 45**

   Copyright © 1982 by CMP Publications, Inc. 111 East Shore Road, Manhasset, N.Y. 11030. Reprinted with permission from *Electronic Engineering Times.*

   Copyright © 1983 by CMP Publications, Inc. 111 East Shore Road, Manhasset, N.Y. 11030. Reprinted with permission from *Electronic Engineering Times.*

10. **Electronic Products, p. 32**

    Copyright © 1983 by Hearst Business Communications, Inc. Reprinted with permission from *Electronic Products,* March 4, 1983 issue.

11. **Graphic Arts Monthly, p. 46**

    Reprinted with permission from *Graphic Arts Monthly & The Printing Industry,* March 1983 issue.

12. **Instruments & Control Systems, p. 25**

    Reprinted with permission from *Instruments & Control Systems*, March 1983 issue.

13. **Machine Design, p. 5**

    Copyright © 1983 by Penton/IPC Publications. Reprinted with permission from *Machine Design*, May 19, 1983 issue.